

# **EXHIBIT A**

---

# EMC Data Manager

---

## Software Reference

EDM Release 4.5.0

P/N 300-113-001-03

Rev C

**EMC<sup>2</sup>**

EMC Corporation, 171 South St., Hopkinton, MA 01748-9103  
*Corporate Headquarters* (508) 435-1000 (800) 424-EMC2  
*Fax* (508) 435-5374 *Service* (800) SVC-4EMC

---

This software/documentation contains proprietary information of EMC Corporation; it is provided under a license agreement that contains restrictions on use and disclosure and is also protected by copyright law. Reverse engineering of the software is prohibited.

The information in this document is subject to change without notice. If you find any problems in the documentation, please report them to us in writing. EMC Corporation does not warrant that this document is error-free.

### **Restricted Rights Legend**

Use, duplication, or disclosure by the Government is subject to restrictions as set forth in subparagraph (c)(1)(ii) of DFARS 252.227-7013, Rights in Technical Data and Computer Software (October 1988).

If this software/documentation is delivered to a U.S. Government Agency not within the Department of Defense, then it is delivered with "Restricted Rights," as defined in FAR 52.227-14, Rights in Data - General, including Alternate III (June 1987).

EMC<sup>2</sup>, EMC, MOSAIC:2000, and Symmetrix are registered trademarks and EMC Enterprise Storage, EMC Enterprise Storage Network, EMC Enterprise Storage Specialist, EMC Storage Logic, The EMC Effect, The EMC Effect Alliance, The Enterprise Storage Company, Celerra, Connectrix, CopyPoint, EDM, E-Infrastructure, FarPoint, InfoMover, PowerPath, SDMS, SRDF, SymmAPI, TimeFinder, Universal Data Tone, and Volume Logix are trademarks of EMC Corporation. Other trademarks are the property of their respective owners.

---

---

# Contents

<b>Preface</b>	Who Should Use This Manual	xxx
	How to Get the Information You Need	xxx
	Online Help and Online Books	xxx
	Hardcopy Documentation	xxxi
	Symmetrix Path and Symmetrix Connect	xxxi
	Clients and Databases	xxxi
	Notes and Cautions	xxxii
	Technical Support	xxxii
	Reader's Comments	xxxii

## **PART I Basic Backup and Restore Concepts**

### **Chapter 1 EDM Basics**

EDM Overview	1-2
EDM Components	1-3
Network Backup and Restore	1-4
EDM Symmetrix Path	1-5

EDM Symmetrix Connect	1-6
Features	1-7
Mirroring	1-7
References	1-8
EDM Hardware	1-9
Disk Subsystem	1-9
Library Units	1-9
External Components	1-10
Remote Diagnostic Modem	1-10
Color System Console	1-10
EDM Software	1-11
Software Components	1-11
EDM Graphical User Interface	1-12
Starting the EDM GUI	1-13
Main Window	1-14
Backup Client Install Wizard	1-15
Backup Configuration Wizard	1-15
Backup Activity Wizard	1-15
Volume Management	1-16
Backup Configuration	1-17
Backup Report	1-18
Restore	1-19
System Monitor Configuration	1-20
HSM Client Installation	1-21
HSM Configuration	1-22

## **Chapter 2 Managing Your EDM System**

Administering the System	2-2
Checking Media	2-3
Verifying Backup Completion	2-3

Reading the Daily Message File	2-4
Saving the EDM Backup Disaster Report	2-5
Check Compaction of Staging and Baseline Volumes	2-6
Running Procedures Automatically via Cron	2-6
Deleting Existing Entries in the crontab File	2-9
Backing Up Server Database Files	2-10

### **Chapter 3 Basic Backup and Restore Concepts**

Client/Server Architecture	3-2
Key Processing Concepts	3-3
Scheduling	3-4
Automatic Scheduling	3-4
Custom Scheduling	3-4
Command Line Scheduling	3-5
Backup Activity Wizard	3-5
Concurrent Work Item Input	3-6
Balanced Scheduling	3-7
Rotation Period	3-7
Load Balancing	3-7
Multiplexed Storage	3-8
Nightly Backup Processing	3-10
Automatic Start	3-10
Client Processing	3-11
Backups of Changing Files	3-11
Storage of Backup Data	3-11
Cataloging of Backup Data	3-11
Restore Processing	3-12

Configuration Options	3-13
Key Configuration Options	3-13
Which Clients to Back Up?	3-14
What Data is Backed Up on Each Client?	3-15
When are Backups Scheduled?	3-15
Where is Backup Data Written?	3-16
Other Configuration Options	3-18
How are Backups Processed?	3-19
Permissions and User Modes	3-19
Catalogs and Backup Savesets	3-20
Reports and Logs	3-21
Reporting in the EDM GUI	3-22
Manual Operations	3-22

## **Chapter 4 Port Control**

Overview	4-2
Understanding Port Control	4-3
Installing and Updating Client(s)	4-3
Restrictions	4-4
What is a Firewall?	4-5
EDM Firewall Assumptions	4-5
Firewall Requirements	4-8
Sample Firewall Configurations	4-9
Port Control Checklist	4-10
Setting Up the Firewall for Port Control	4-12
Enabling Port Control on the EDM	4-12
Portservices Files	4-14
localhost File	4-17

Installing Port Control on the EDM Client(s)	4-18
NOS Client Access Window	4-18
UNIX Client Install Method Window	4-18
Port Control Window	4-18
Making Changes to Port Control	4-20
To Change System Monitoring for Ping Errors	4-20
To Change the Default Port Range	4-20
An Alternative Method	4-20
To Enable Port Control with eb_server_config:	4-21
To Turn Off Portmapper on the EDM	4-22
An Alternative Method	4-22
To Turn Off Portmapper on Client(s)	4-22
To Set Portmapper Off for New Clients	4-22
To Disable Port Control for a Single Client	4-23
To Disable Port Control on the EDM and All of Its Clients	4-23

## **Chapter 5 How Backup and Restore Work**

How Backup Works	5-2
Monitoring Active Backups	5-4
Start of Overall Backup	5-4
Client Access	5-4
Work Item Specification	5-5
Automatic Scheduling	5-5
Backup Activity Wizard	5-5
Client/Server Processing Methods	5-6
Standard Client Processing	5-6
High-Speed Client Processing	5-6
Server Processing	5-7
Filesystem Backup	5-7
Client Scans Filesystem	5-7



Database Backup	5-8
ACL Support	5-8
Backing Up Files with ACLs	5-8
Restoring Files with ACLs	5-9
Cross-Client Restore	5-9
Client ACL Commands	5-10
Client Pacing	5-10
Server Processes Attributes and Data	5-12
Catalog Processing	5-13
Server Database Update	5-14
Report and Log File Generation	5-14
Backup Completion Reports	5-14
Backup Failure Reports	5-15
Backup and Restore Logs	5-15
How Restore Works	5-16
Media Management	5-19
Expiring Backups	5-20
Deallocating Baseline Volumes	5-20

## **Chapter 6 Database Backup and Restore**

Varieties of Database Backup	6-2
Database Backup Clients	6-3
EDM Symmetrix Path	6-3
EDM Symmetrix Connect	6-4
Various Database Backup Clients	6-6
Oracle Backup Client	6-6
EMC Backint Client for SAP R/3 Oracle Databases	6-6
Other UNIX Database Backup Clients	6-6
Microsoft Database Backup Clients	6-7

Database Network Backup Overview	6-7
Multiple Streams	6-8
Server-side Processing	6-8
Restores	6-9
User Interfaces	6-9
Configuring Backups to an Alternate Network	6-9
Database Pre-Discovery	6-10
EDM Symmetrix Path Overview	6-12
EDM Symmetrix Connect Overview	6-13
Applications	6-13
User Interfaces	6-13
Documentation	6-14
Raw Device Backups	6-14
Configurations	6-15
Mirrored Configurations	6-16
Non-Mirrored Configuration	6-17
EDM's Legacy "Offline" Database Backup Feature	6-17

## **Chapter 7 Basic Volume Management Concepts**

Volume Management Overview	7-2
EDM Library Unit Manager	7-3
Volume Manager	7-3
Volume Catalog	7-4
Volume Life Cycle	7-4
Uncataloged	7-5
Unlabeled	7-6
Foreign	7-6
Expired	7-6
Erasing	7-7

Unverified Volume	7-7
How Volumes are Allocated	7-8
When a Volume is Allocated	7-8
Library Managers	7-12
Library Manager Configuration	7-12
Robotic Library Units	7-13
Offline and Offsite Library Managers	7-13
Ejecting a Volume	7-14

## **Chapter 8 How Volume Management Works**

rvmoper UNIX Group	8-2
Volume Management Processes	8-2
Volume Manager	8-3
Library Managers	8-4
Notify Daemon	8-4
Volume Management Startup	8-4
Manually Stopping and Restarting the vmdaemon	8-5
Using edmproc -restart	8-5
If an Error Occurs While Using edmproc -restart	8-7
Library Unit Operations	8-8
Inserting Media into Library Units	8-8
Importing a Volume	8-9
Inserting Cleaning Cartridges into Library Units	8-10
When Drives are Busy	8-11
Mounting and Dismounting Volumes	8-12
Mounting Volumes	8-12
Dismounting Volumes	8-13

Ejecting Media from a Library Unit	8-14
Ejecting Media Through the EDM GUI	8-15
Ejecting Media at the CLI	8-15
Drive Scheduling and Preemption	8-16
Drive Preemption	8-16
Verifying Priority in the Queue	8-16
LM_MAX_RESIDENT_TIME and	
LM_MIN_RESIDENT_TIME	8-17
Simultaneous Backup Example	8-17
Library Unit Inventories	8-18
Inventory Tables	8-19
How Inventories are Done	8-20
Delta Inventory	8-21
Barcode Inventories	8-21
Cleaning Tape Drives	8-22
Volume Allocation and Deallocation	8-23
How Volumes are Allocated	8-23
Volume Allocation Request	8-24
Mount Request	8-24
Volume Use	8-25
Duplicate Volume Sequence Numbers	8-26
When Volumes are Deallocated	8-27

## **Chapter 9 Media Duplication**

The Media Duplication Process	9-2
Media Duplication Commands	9-3
Starting Duplication	9-3
Preparing a Backup Volume for Duplication	9-4
Selecting Append Mode or New Mode	9-5
Append Mode	9-5
New Mode	9-6

Configuring a Trail for Duplication	9-8
Initiating Manual Duplication at the CLI	9-8
Determining Duplicates of an Original Volume	9-10
Setting the Maximum Number of Concurrent Duplications	9-11
Configuring Concurrent Duplications in the EDM GUI	9-12
Configuring Concurrent Duplications at the CLI	9-12
Verifying the Status of a Duplication	9-13
If a Duplication Was Scheduled for an Offline Volume	9-14
Manually Disabling and Re-enabling Duplication	9-14
Disabling Duplication	9-15
Re-enabling Duplication	9-15
Pausing, Resuming, Canceling, or Removing a Duplication	9-16
Pausing Duplication	9-16
Resuming Duplication	9-17
Canceling Duplication	9-17
Removing a Failed Duplication from the Queue	9-18
If a Duplication Fails	9-19
Rescheduling a Failed Duplication	9-19
Rescheduling a Failed Duplication Through the GUI	9-20
Rescheduling a Failed Duplication at the CLI	9-21
Rescheduling Duplication of an Offline Original Volume	9-22
Rescheduling Duplication of a Single Volume for Archival Purposes	9-23

Restoring from Backup or Duplicate	9-24
If an Original Volume is Defective	9-25
Viewing Reports on Duplications	9-25
ebreport media Report	9-25
ebreport duplicate Report	9-27
Importing a Duplicate Volume	9-29
Importing a Duplicate Volume before the Original	9-29
Rejecting a Mount Request	9-29
Expiring a Duplicate Volume	9-30
Viewing Duplicate Expiration Dates	9-31

## **Chapter 10 Magnetic Disk Concepts**

Expiration of Backups and Catalogs	10-2
Choosing Expiration Periods	10-3
Running Expiration	10-4
Filesystem Cleanup Script	10-4
Magnetic Disk Capacity	10-5
Rotation and Keep Catalog Periods	10-5
Calculating Actual Daily File Changes	10-7
Managing Disk Space	10-10
Distributing Catalogs	10-10
Reclaiming Magnetic Disk Space	10-13
Manually Expire Unneeded Catalogs	10-13
Other Options	10-14
Changing the Automatic Cleanup Script Defaults	10-15

## **PART II Hierarchical Storage Management (HSM)**

### **Chapter 11 Basic HSM Concepts**

- When Files Stage In and Out 11-2
- Filesystem Configuration and Maintenance 11-3
  - Staging Templates and Staging Trails 11-4
    - Deciding How Many Staging Templates to Create 11-5
  - Bitfiles and Client Stores 11-6
    - Deciding How Many Client Stores to Create 11-7
    - Preventing Redundant Backup of EDM Migration Client Data 11-9
  - Full Filesystems 11-10
  - Watermarks 11-10
    - Disk Utilization Zones 11-13
    - Sample Watermarks 11-15
  - Configuration Issues 11-18
    - Filesystem Limits 11-18
    - Stage-to-Tape Considerations 11-19
    - Self-Describing Media 11-22
  - Periodic Staging and Filesystem Delay 11-22
- File Control Properties 11-23
  - Listing and Changing File Control Properties 11-24
- Compaction of Staging Media 11-27
  - Administering Compaction 11-28
- Compacting Baseline Media 11-29
- Migration Reports 11-30
  - emfsreport and the Working Set 11-31
- Baseline Backup 11-33

Restaging Data	11-33
Backup Completeness	11-34

## **Chapter 12 How Migration Works**

What Happens When You Enable Migration	12-2
Migration Configuration Database	12-2
How Stage-Out Works	12-4
The File Monitor Daemon (emfmd)	12-4
The Master Staging Daemon (emmasterd)	12-5
Candidate List Generation	12-5
What Happens When a File is Staged Out	12-6
How Stage-In Works	12-7
How the User-Level Commands Work	12-7
How the Network Migration Server Works	12-8
The EDM Migration Protocol	12-8
Network Migration Server Daemons	12-9
Network Migration Server Database	12-10
Client Stores	12-12
Bitfiles	12-13
How Compaction Works	12-14
Compaction Goals	12-14
Example	12-15
The Compaction Process	12-16
How Long Compaction Takes	12-18
Baseline Compaction	12-18
Deallocation and Reuse	12-18
Active Baseline Volumes	12-19
Recovering from Site Disasters	12-19



**Chapter 13 HSM Command-line Tasks**

- HSM Commands 13-2
- Test Staging 13-2
- Set Up Periodic Staging 13-2
- Tuning for Staging to Tape 13-3
  - Stage-to-Tape Tuning Parameters 13-3
  - Stage-to-Tape Tuning Procedure 13-4
- Coordinating Automatic Procedures 13-4
  - Client (Periodic Stage Out) 13-5
  - Server (emvck) 13-6
  - Server (Periodic Stage Out) 13-6
  - Server (Compaction) 13-6
  - Server (Baseline Backup) 13-7
  - Server and Client (Backup) 13-7
  - Server (Backup Database) 13-7
  - Rotating Error Logs 13-8
- Working with Individual Files 13-9
  - Staging Out Files 13-9
  - Staging In a Set of Files 13-9
  - Tagging a Set of Files for Future Stage Out 13-10
  - Locking a File on Magnetic Disk 13-11
- Checking the Staging Configuration 13-12
- Copying and Moving Data 13-13
  - Migrating Data from One Staging Trail to Another 13-13
  - Copying Files from One Filesystem to Another 13-14

Moving and Copying Files Between HSM Systems	13-14
Copying Files to Another HSM System (No Media)	13-15
Moving Files to Another EDM (Media Included)	13-15
Moving Files Between HSM Clients (Store Included)	13-16
Restoring a Staged Out File That Has Been Deleted	13-17
Copying Files to a Non-EDM System	13-17
Monitoring Storage Space and File Sizes	13-18
Maintaining Non-Stageable Filesystems	13-19
Managing Your Magnetic Disks	13-19
Run emfsreport	13-20
Choose the Desired Working Set in Days	13-21
Determine Additional Magnetic Disk Space Required	13-22
Reconfigure Magnetic Disk or Purchase More Disks	13-23
Populating Filesystems	13-23
Disabling Filesystem Staging	13-24
Temporarily Disabling Periodic Staging	13-25
Permanently Disabling Periodic and Demand Staging	13-25
Moving Staged Out Files to Another Filesystem	13-27
Compacting Staging Media	13-28
Administering Compaction	13-29
Compacting Baseline Media	13-30

- Clearing Incomplete Bitfiles 13-31
- Moving a Store to Another EDM Server 13-31
- Gathering Migration Store Statistics 13-32
- Checking a Network Client's Staging Configuration 13-33
- Troubleshooting HSM 13-33
- Restoring a Lost or Damaged Staging Volume 13-34
  - Restoring a Lost or Damaged Staging Volume (No Baseline Backup) 13-34
  - Restoring a Lost or Damaged Staging Volume (Baseline Backup is Enabled) 13-35
- Restoring a Lost or Damaged Staging Trail 13-35
- Restoring a Lost or Damaged Filesystem 13-36

## **PART III Logs and Reports**

### **Chapter 14 Start of Backup and Related Processing**

- Backup Processing 14-2
  - Backup Activity Wizard 14-3
  - Automatic Nightly Processing 14-3
    - Automatic Scheduling 14-6
    - Custom Scheduling 14-6
  - Command Line Processing 14-7
- Catalog Processing 14-9

### **Chapter 15 Message Logging**

- Message Logging Features 15-2
- Syslog Message Files 15-2

Circular Log Files	15-4
Log File Rotation and Archival	15-4
Log Message Format	15-5
Default syslog Configuration File	15-6

## **Chapter 16 Backup Reports and Log Files**

Report and Log Usage	16-2
Executing Reports from the EDM GUI	16-3
Report and Log Summaries	16-4
Backup Reports	16-5
Backup Media Reports	16-10
Backup Duplicate Reports	16-12
Duplicate Command Options	16-13
Sample Backup Duplicate Report	16-14
Backup History Reports	16-15
History Command Options	16-16
Sample Backup History Report	16-18
Backup Disaster Reports	16-19
Backup Baseline Reports	16-27
Backup Completion Reports	16-29
Backup Failure Reports	16-31
Backup Coverage Reports	16-32
Volume Reports	16-33

Log Files	16-35
Server Log Files	16-35
Local Client Log Files	16-37
Remote Client Log Files	16-37
Other Logs	16-37

## **PART IV Command Line Interfaces**

### **Chapter 17 Configuring Library Managers**

Using the lmconfig Utility	17-2
Listing Library Managers	17-3
Library Manager Name	17-3
SCSI Address	17-4
Installing Device Drivers	17-4
Updating Device Drivers	17-6
Removing Device Drivers	17-7
Configuring a Library Manager	17-8
Preparing for Configuration	17-8
Running lmconfig	17-9
If All Library Units Are Configured	17-10
If Media Is Found In Any Drive	17-10
The Configuration Process	17-11
Selecting the Cleaner Barcode Default	17-13
Viewing Log Files	17-13
Completing lmconfig	17-14
Deconfiguring a Library Manager	17-14

If You Have Trouble Configuring a Library Unit	17-17
If a Problem Occurs While Configuring Multiple Library Units	17-17

## **Chapter 18 Man Page Listing**

Backup and Restore Man Pages	18-2
Volume Management Man Pages	18-6
HSM Man Pages	18-9

## **PART V Disaster Recovery**

### **Chapter 19 Being Prepared for a System Disaster**

Safeguarding Your Backup Media	19-2
Running and Saving Reports	19-3
MINIMAL Disaster Report	19-4
FULL Disaster Report	19-4
Redundant Backup Coverage	19-5
Configure Alternate Media Sets (Trailsets)	19-5
Media Duplication	19-5

### **Chapter 20 Recovering a Server from a Disk Failure**

Steps to Restore a Server	20-4
Stop All Activity on the Server	20-5
Disable Activity	20-5
Reinstall Hardware and Software as Needed	20-6

Temporarily Reconfigure the Server	20-11
lmconfig	20-11
eb_server_config	20-13
Restore LOCAL_DATABASE Files	20-15
Reconfigure Library Units	20-23
Restore Catalogs and Backup Information Created After LOCAL_DATABASE Backup	20-23
Stop Backups	20-23
Run a Full Inventory and Import Uncataloged Volumes	20-24
Import Backup Catalogs and List Volumes	20-24
Restore Catalogs and Backup Information	20-26
Restore Data Created Before LOCAL_DATABASE Backup	20-28
Reenable crontab Entries	20-30
Restore Past Catalogs	20-30
Restore Missing Catalogs	20-32

## **Chapter 21 Recovering a UNIX Client from Disk Failure**

Recovering a Client	21-2
Beginning Steps	21-2
For HSM Clients Only	21-3
For Backup and HSM Clients	21-4
For HSM Clients	21-5
For Both Backup and HSM Clients	21-6

## **PART VI Appendices**

### **Appendix A Directory Structure**

Backup Server Directory Structure	A-2
/usr/epoch	A-2
/usr/epoch/EB	A-4
Bin Directory	A-5
Catalogs Directory	A-5
Client Directory	A-6
Config Directory	A-8
Db Directory	A-9
Locks Directory	A-10
Log Directory	A-10
Preconfig Directory	A-10
Server Man Directory	A-10
Table of Backup Server Directories and Files	A-11
Backup Client Directory Structure	A-13
Client Home Directory	A-13
Security Issues	A-15
Additional Client Software	A-15
Client EB_DB Directory	A-16
Client bin Directory	A-16
Client man Directory	A-16
Table of Backup Client Directories and Files	A-17
Volume Management Directory Structure	A-18
Library Manager Subdirectories	A-18
Volume Manager Files	A-20



HSM Directory Structure	A-21
HSM Configuration Database	A-21
Network HSM Server Database	A-22
Client Stores	A-23
Bitfiles	A-25

## **Appendix B EDM Backup Configuration File**

General Coding Rules	B-2
Checking Your Changes	B-3
When Changes Take Effect	B-4
Summary of Fields	B-4
Server Fields	B-16
ebserver	B-17
Client Backup Username	B-17
Backup Administrator Usernames	B-17
Authorized Backup List	B-18
Temporarily Disabling Backups for a Client	B-19
Authorized Recovery List	B-19
Authorized Cross-Recovery List	B-20
Disabling All Cross-Restores	B-22
Recovery Administrator List	B-23
Disabling Administrator-Level Restores	B-24
Maximum Simultaneous Clients	B-24
Use At Most n media-type Trails Concurrently	B-24
Limit Throughput To: nnn Per time	B-25
Specifying No Throughput Limit	B-27
Maximum Server backups.log File Size	B-27
Maximum Server recoveries.log File Size	B-28
Maximum Client backups.log File Size	B-28

Maximum Client recoveries.log File Size	B-29
Catalog Threshold to Force Level 0 Backup	B-29
Work Group Fields	B-31
Work Group Name	B-31
Filesystem Work Item Fields	B-32
When You Change a Work Item or a Filesystem	B-35
When You Stop Using a Work Item	B-35
Work Item Name	B-36
Client Name	B-37
Filespec to Back Up	B-37
Using the Block Form of the Syntax (Filespec Statement)	B-37
Backing Up the Local Client	B-38
Baseline Filespec	B-38
Migration Backup Tag	B-38
Exclusion Tag	B-40
Connection Via	B-40
Setting up the Network	B-41
Configuring the Work Items	B-41
Priority	B-41
Do Not Load Balance	B-43
Completeness	B-44
Level Map	B-45
Access Time Preservation	B-48
Maximum Files not Backed Up Before Forcing Full Backup	B-49
Database Work Item Fields	B-50
Database Work Item Name	B-51
Client Name	B-52
Filespec	B-52

Partition Spec	B-52
Database Type	B-52
Database Server	B-52
Database Name	B-52
Type	B-53
Exclusion Tag	B-53
Inclusion Tag	B-53
Access Time Preservation	B-53
Priority	B-54
Do Not Load Balance	B-54
Backup Client Initialization Command	B-54
Initialization Timeout	B-55
Backup Client Cleanup Command	B-55
Cleanup Timeout	B-55
Buffer Sizes	B-55
Backup Start Time	B-56
Level Map	B-56
PC Work Item Fields	B-56
Work Item Name	B-57
Connection Method	B-57
Filespec	B-57
Netware Username	B-58
Netware Encrypted Password	B-58
Netware Client TSA	B-58
Netware Client Target	B-58
Exclusion Tag	B-58
Backup Trailsets	B-58
Backup Trailset Name	B-60

Use Trail	B-60
Trail Name	B-61
Media Type	B-61
Backup Level	B-61
Maximum Number of Concurrent Clients	B-62
Use Level B1 / B2 For Baseline Backups	B-64
Keep Backups	B-65
Keep Duplicates	B-67
Keep Backup Catalogs	B-68
Keep Saveset Records	B-68
Backup Catalog Delta Level	B-69
Backup Template Fields	B-70
Backup Template Name	B-71
Work Group List	B-73
Begin Trailset Rotations	B-73
Rotation Period	B-74
Primary and Alternate Trailsets	B-75
Logging Level	B-77
Server Log File	B-78
Backup Completion Script	B-79
Backup Failure Script	B-80
Do All Baseline Backups Before Normal Backups	B-80
Recreate Baseline if Needed	B-81
Schedule	B-82
Standard vs. Full-During-Weekends Rotations	B-83
Specifying Backup Shifts	B-84
Scheduling Custom Backups (Level n on Days...)	B-84
Startup Parameters	B-86

**Appendix C Volume Management Configuration Files**

- Volume Manager Configuration File *C-2*
- Library Manager Configuration Files *C-9*
  - Library Manager Naming Convention *C-9*
  - Library Manager Subdirectories *C-10*
  - Library Manager Configuration Parameters *C-11*
    - The LM\_INLET\_IGNORE\_ON\_OPEN  
Parameter *C-19*
  - Offline and Offsite Library Managers *C-20*
    - Offline Library Manager *C-20*
    - Offsite Library Manager *C-21*

**Appendix D findxcpio Directives**

- Work Item Directive *D-2*
- Logical Operators *D-2*
- Macros *D-3*
  - Compound Macros *D-5*
- Syntax to do Back Up *D-7*
- Syntax to not Back Up *D-8*
- Evaluation Shortcuts *D-11*
  - False and *D-11*
  - True or *D-12*

**Glossary****Index**

---

## Preface

The *EMC Data Manager Software Reference* manual is the core document that is provided to customers who purchase either an EMC Data Manager (EDM) or EDM with the Hierarchical Storage Management (HSM) Option.

This manual provides comprehensive reference information about Backup, Restore, Volume Management, and optional HSM software. This manual introduces basic EDM software concepts and describes the programs, reports, and log files that you use in conjunction with the EDM interface. It also provides overall disaster recovery instructions for the EDM server.

This manual describes the following:

- Basic Backup and Restore Concepts
- Volume Management and Duplication Concepts
- Basic HSM Concepts
- Logs and Reports
- Command Line Interfaces
- Disaster Recovery
- Configuration Files

---

## Who Should Use This Manual

This manual is intended for system administrators who are responsible for administering and operating an EMC Data Manager (EDM) or EDM with HSM Option. You should be familiar with UNIX system administration, understand your network environment, and understand the requirements of the various groups that you serve.

---

## How to Get the Information You Need

Information about the EMC Data Manager is available through the graphical user interface (GUI) online help and online books, as well as hardcopy documentation, as described in the following sections.

---

## Online Help and Online Books

The EDM graphical user interface Help facility provides detailed information and procedures for managing the EDM server, its clients, and library units.

To access online help, log in to the EDM. Enter **edm** to display the EDM Main window, then click on the Help button. The Help facility includes context-sensitive help for input fields, window areas, options, and buttons. Online Help also includes step-by-step instructions for tasks that you perform to manage and monitor the EDM server.

Also available are online versions of selected manuals. To access these online books, select Help in the EDM Main window menu bar, then Online Books.

---

## Hardcopy Documentation

The following documentation is available for use with EMC Data Manager (EDM) or EDM with HSM Option.

- *EDM Software Reference* (P/N 300-113-001)
- *EDM Software Release Notes* (P/N 300-113-004)
- *EDM Storage Devices* (P/N 300-113-002)
- *EDM Server Error Messages* (P/N 300-113-014)

---

## Symmetrix Path and Symmetrix Connect

The *EDM Symmetrix Path User Guide* (P/N 300-113-007) provides information for configuring, backing up, and restoring data using EDM Symmetrix Path.

The *EDM Symmetrix Connect User Guide* (P/N 300-113-005), the *EDM Symmetrix Connect Quick Reference Card* (P/N 300-113-006), and the *Symmetrix Connect Checklist* (P/N 300-113-011) provide information for configuring, backing up, and restoring data using EDM Symmetrix Connect.

---

## Clients and Databases

The following EMC Data Manager supplemental guides provide installation and configuration instructions for setting up the EDM backup clients. Release notes are also available.

- *NetWare Backup Client* (P/N 300-114-001)
- *OS/2 Backup Client* (P/N 300-118-001)
- *Windows NT Backup Client* (P/N 300-119-001)
- *OpenVMS Backup Client* (P/N 300-122-001)
- *Oracle Backup Client* (P/N 300-115-001)
- *Sybase Backup Client* (P/N 300-116-001)
- *Informix Backup Client* (P/N 300-117-001)
- *EMC Backint for SAP R/3 System* (P/N 300-120-001)
- *Windows NT SQL Server Backup* (P/N 300-121-001)
- *Windows NT Oracle Backup Client* (P/N 300-115-003)



- *Windows NT Exchange Backup Client* (P/N 300-119-003)
- *Windows NT Lotus Notes Backup Client* (P/N 300-119-005)

---

## Notes and Cautions

Notes provide clarification or additional important information.

**Note:** It calls your attention to an operating procedure, practice, condition, or similar situation which is important to highlight.

Cautions are used to indicate the presence of a hazard.

**CAUTION:** It calls your attention to an operating procedure, practice, condition, or similar situation. Failure to observe a caution can result in minor personal injury, damage to a program, device, system, or loss of data.

---

## Technical Support

If you need technical assistance with the EDM system, contact the EMC support hot line at:

**1 800 SVC-4EMC (1 800 782-4362)**

If you are located outside the United States or Canada, contact the nearest EMC office for assistance.

---

## Reader's Comments

We welcome your comments on this documentation.

Send e-mail to:

**`doc_comments@mil.emc.com`**

Please include the part number and title of the manual.

---

---

**Part I**  
**Basic Backup**  
**and Restore**  
**Concepts**

---



---

# 1 EDM Basics

EMC Data Manager (EDM) is an integrated solution for unattended backup of data over your network, EDM Symmetrix Path, or EDM Symmetrix Connect.

The EMC Data Manager is a full-featured, client/server backup system that offers fast and reliable backup processing with minimal operator intervention.

With its graphical user interface (GUI) and built-in intelligent scheduling, EDM Backup software fully manages the complete backup process to ensure maximum protection of all data in a client/server environment.

This chapter describes the following topics:

- EDM Overview
- EDM Hardware
- EDM Software
- EDM Graphical User Interface

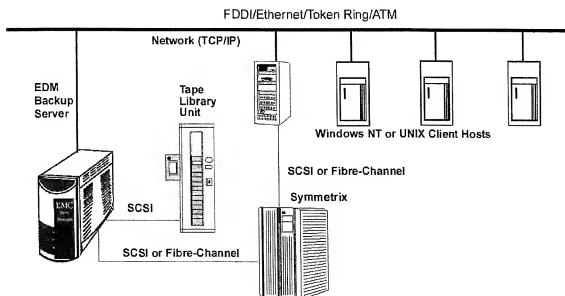
## EDM Overview

EMC Data Manager backs up data from magnetic disk to the magnetic tape, backing up both databases, tablespaces, and filesystems. EDM writes the backup data to tape cartridges within one or more tape library units attached to the EDM. In so doing, EDM gives you the ability to perform restores of data lost by physical disk failure and restores of data lost by removal (logically) from the disk.

EDM backs up data from both Symmetrix-attached computers and from networked computers. For Symmetrix-attached computers, EDM's high-performance Symmetrix Connect and Symmetrix Path options offer rapid transfer of large amounts of backup data over Fibre-Channel or SCSI cabling between the Symmetrix and EDM. For computers that are not attached to a Symmetrix, EDM transfers backup data over the network.

Figure 1-1

EDM Symmetrix and Network Backup Environment



With its Hierarchical Storage Management (HSM) option, EDM migrates less-used filesystem data from magnetic disks to optical disks located in attached optical library units. One benefit of HSM is to reduce full backup loads.

The EDM Backup software maintains catalogs of the backups on disk storage, and restores filesystem and database data from magnetic tape to the client.

The EDM software also controls the operation of robotics, drives, and cartridges located inside tape library units. State-of-the-art product features include a multi-threaded library manager, media duplication capability, and AT&T StorLink support.

A graphical user interface enables you to configure and manage tape operations, and network backups and restores of backed up files and databases.

---

## **EDM Components**

The EDM system includes both the hardware and software components that are needed to backup and restore data. The EDM hardware consists of:

- an EDM cabinet containing: a server unit, magnetic disks for online catalogs, and optional internal tape library unit
- a color system console, keyboard, mouse, and modem
- attached external tape library units (DLT, 8mm) and/or optical library units (EO, WORM)


EDM software components include a Sun Solaris operating system and EDM Backup software (which includes system monitoring support software). The software is installed by EMC and vendor factory and/or EMC Customer Service personnel. Optionally included are Sun ATM, FDDI, Ethernet, and/or Token Ring card device driver software, and EMC Hierarchical Storage Management (HSM) software.

## Network Backup and Restore

Over the network, EDM can back up from a broad selection of UNIX and Windows NT platforms and PC clients acting as file servers, database management systems, and application workstations. Computer platforms in the EDM network environment include the EDM itself and the other hosts that you configure as network backup clients. (Refer to the *EMC Data Manager Software Release Notes* for a current list of supported clients for network and direct connect backup.) Data from each host is streamed over the TCP/IP network to the EDM, which writes the data to tape library units (TLUs) that are attached to it through SCSI connections. Figure 1-2 illustrates an EDM network environment.

Figure 1-2

EDM Network Backup and Restore Environment



The EDM software maintains catalogs of the backups to enable easy interactive restoring whenever necessary. The backup catalogs are stored online on the disk subsystem in or external to the EDM cabinet.

EDM performs online, full and incremental backups of filesystems. Full backups copy the entire filesystem, while incrementals just copy those files that have been added or changed since the previous backup. (Incremental backups also note down any files that have been deleted.) EDM can back up databases online or offline, as applicable to particular database systems.

As EDM system administrator, you can automate filesystem backups and you can also initiate backups and restores on demand from the EDM. Also, you can optionally configure the restore feature to enable some or all of your workstation and fileserver users to restore their own files and filesystems without your assistance.

In most cases, database backups can be either client-initiated or EDM-initiated. With client-initiated backups, the database administrators manage their own database backups and restores from the database server (the client system to the EDM). With EDM-initiated backups, you can configure automated schedules as well as initiate backups on demand. In a few cases, database backups can only be EDM-initiated, and in a few cases, only client-initiated. Similarly, in many cases, database restores can either be EDM-initiated or client-initiated. But in some cases, database restores can be EDM-initiated only, and in other cases, client-initiated only.

---

## **EDM Symmetrix Path**

With EDM Symmetrix Path, you can back up and restore many of the same types of databases and filesystems that you can over the network, but the data is transferred through a Symmetrix instead of over the network. (Refer to the *EMC Data Manager Software Release Notes* for a current list of supported client platforms.) The EDM Symmetrix Path option backs up large databases and filesystems through the Symmetrix cable connections, thereby bypassing the network's limited speed and bandwidth.



With this methodology, the EDM server and a client system are both cabled to a Symmetrix and the Symmetrix itself acts as the network. A few small devices on the Symmetrix are designated as transport paths for configuration purposes, while the actual data transport is generally handled by cache memory. Various device configurations and host mapping combinations are possible to provide different performance characteristics and degrees of flexibility.

Symmetrix Path can backup data residing on the locally attached disks of a Symmetrix-attached computer, as well as data on Symmetrix disks. Refer to the *EMC Data Manager Symmetrix Path User Guide* for more information about this feature.

---

## EDM Symmetrix Connect

The EMC Data Manager Symmetrix Connect feature permits a high speed online or offline backup and restore of a client's very large UNIX filesystem, Windows NT, Oracle7, Oracle8, or Oracle8i database (or a copy) that resides on one or more Symmetrix systems.

EDM Symmetrix Connect backs up data residing on Symmetrix disks, either on so-called mirrored volumes — TimeFinder Business Continuance Volumes (BCVs) or Symmetrix Remote Data Facility (SRDF) target (R2) volumes — or on unmirrored volumes — TimeFinder standard devices (STDs) or SRDF source (R1) volumes.

Symmetrix Connect provides EDM-specific interfaces to back up database and filesystem data on UNIX and Windows NT clients. For UNIX systems (IBM, Compaq, HP, Sequent, and Sun), the EDM Oracle Application Interface backs up Oracle7/8/8i databases and filesystems. On Windows NT clients, various EDM-specific interfaces enable backup of filesystems, Oracle Backup, SQL Server, and Exchange.

Symmetrix Connect also backs up UNIX Oracle databases through two standard interfaces. One is Oracle8's Recovery Manager (RMAN) using its Proxy Copy feature. The other is the SAP R/3 System SAPDBA utility through EMC's SAP-certified interface client: EMC Backint.

---

## Features

Symmetrix Connect offers these features:

- high speed backup throughput to tape media
- Oracle database and tablespace-level online and offline backup capability, plus data file-level backup if using RMAN Proxy Copy
- data file-level restore (plus disaster recovery) capability
- minimal impact on your local area network
- minimal impact on the client (whose database is being backed up)
- UNIX filesystem backup for all supported clients
- Oracle filesystem backup for all supported clients
- Windows NT support for Oracle Backup, Exchange Server Backup, SQL Server Backup, and filesystems
- SAP R/3 Symmetrix Connect support for Oracle databases on UNIX clients database running EMC Backint
- multiple simultaneous database support for UNIX clients
- PowerPath interoperability
- GUI client configuration and backup capability
- enhanced GUI System Monitoring Support features

---

## Mirroring

If you are using a mirroring facility, Symmetrix Connect backups take advantage of three Symmetrix features to provide high performance database backup with minimal impact on your database server (host):

- Symmetrix Remote Data Facility (SRDF)
- Symmetrix TimeFinder (or SMMF™)
- Symmetrix Remote BCV in an SRDF Configuration

In addition, the multipath disk access feature is Symmetrix's ability to allow multiple (local and remote) hosts to read and write data to and from the same disk at the same time.

**Note:** Hierarchical Storage Management (HSM) is not supported with EMC Data Manager Symmetrix Connect software.

---

## References

For more information on Symmetrix Connect, refer to:

- The *EMC Data Manager Symmetrix Connect User Guide* for information on using this feature with the EDM Oracle Application Interface for UNIX clients and Filesystem Application for UNIX clients.
- The following guides for information on Symmetrix Connect with the EDM-specific interfaces to NT filesystems and databases:
  - *EMC Data Manager Windows NT Backup Client*
  - *EMC Data Manager Windows NT Oracle Backup Client*
  - *EMC Data Manager Windows NT SQL Server Backup Client*
  - *EMC Data Manager Windows NT Exchange Backup Client*.
- The *EMC Data Manager Oracle Backup Client* guide for information on using Symmetrix Connect with RMAN Proxy Copy (on UNIX clients).
- The *EMC Data Manager EMC Backupint* guide for information on using Symmetrix Connect with the SAP R/3 System's SAP Tools (on both UNIX and Windows NT clients).

---

## EDM Hardware

You can purchase an EDM in several configurations (models) and with various system upgrades.

The server unit can be configured with multiple SBus cards; their use depends on your SCSI peripheral and networking requirements. SBus board options include Ultra-SCSI (differential), EDM-Fibre (Fibre-Channel), ATM, Quad Ethernet, Fast Ethernet, Token Ring, and FDDI.

---

## Disk Subsystem

In addition to a Power Distribution Unit (PDU) and a server unit, most EDM cabinets contain a disk subsystem that stores catalogs and contains operating system and application software.

If you have a SPARCserver 1001E or Ultra Enterprise 4000 (or 4500) unit, a Sun disk subsystem is provided for the storage of backup catalogs, EDM Backup software, Solstice, and other files. An Ultra Enterprise 3000 (or 3500) system contains six or more internal disk drives that function in a similar manner.

The disk subsystem catalogs all files that were backed up by EDM. The disk subsystem is required in network-only (non-Symmetrix) backup environments.

For concepts on the use capacities of the various catalog disk subsystems, refer to "Magnetic Disk Capacity" on page 10-5.

---

## Library Units

EDMs are equipped with tape library units (DLT, DTF, 111TC, 8mm) and/or optical library units (EO or WORM) that provide secondary online storage and perform unattended backup of user data. Each library unit holds a robotics or picker system, one or more drives (tape or optical), and media that enables you to store from tens of gigabytes to tens of terabytes of data.

Refer to *EMC Data Manager Storage Devices* for more information about supported library units.

---

## External Components

In addition to the EDM cabinet(s) and any external library units, your EDM comes with the following hardware:

- Remote Diagnostic Modem
- Color System Console

Installation of Symmetrix systems is handled separately.

---

### Remote Diagnostic Modem

A remote diagnostic modem (RDM), external to the cabinet and equipped with RS-232 cable, enables dial-in. The modem must be connected to a dedicated telephone line.

The modem and the Remote System Monitoring (RSM) function enable the EDM to notify the EMC Customer Service Database about general system information and any problems with the EDM system. The modem and RSM also enable remote-user dial-in by EMC personnel to query the EDM.

The EMC Customer Service Database is the service management software that Customer Engineers, Customer Service Technicians, and Product Support Engineers use to log service activity, track field service inventory, download files, and to look up other important customer information.

---

### Color System Console

The Sun system console has a color, high-resolution, bit-mapped display and includes a keyboard, and a three-button optical mouse.

The color system console uses the Common Desktop Environment (CDE) for displaying the EDM GUI. It enables GUI-based system administration of the EDM as well as access to applications that run on network accessible hosts.

---

## EDM Software

The EDM Backup software with client/server architecture supports Symmetrix Connect and Symmetrix Path as well as standard network backups. The server software is located on the EDM unit. Client software is located on the EDM and on the other hosts within the network architecture.

For more information, refer to "Client/Server Architecture" on page 3-2.

---

## Software Components

EDM software consists of the following components:

- Backup — provides centralized backup and restore services for the server and clients on the network. The EDM's backup server software runs on the EDM server. The client software runs on the server as a local client and on each networked (remote) client. To understand network backup and restore, review Chapter 3 "Basic Backup and Restore Concepts," Chapter 5 "How Backup and Restore Work," and Chapter 6 "Database Backup and Restore."
- Volume Management — provides integrated media and library management for the EDM server. The volume management software keeps track of all media that is known to the server, whether it is in a library unit or an offline or offsite location. For an understanding of volume management, review Chapter 7 "Basic Volume Management Concepts," Chapter 8 "How Volume Management Works," and Chapter 9 "Media Duplication."
- HSM (Hierarchical Storage Management) — is an option that is available with EDM network backups. This option provides HSM for the EDM system, both for the local server and for networked clients. For an understanding of HSM, review Part II, "Hierarchical Storage Management."

- **System Monitoring Support** — enables Customer Service and/or system administrators to configure and receive notification of serious system problems that prevent successful completion of the backups or related issues and system status. When configured with RSM software, notifications can be sent to EMC's Customer Service Database system and to designated email recipients. SNMP traps can be generated and Tivoli event messages issued. Use the System Monitor Configuration GUI to configure it, and online Help to learn more about its features.
- **Client** — is available for a wide range of PC, Windows NT, and database clients.

---

## EDM Graphical User Interface

For network applications, use the EMC Data Manager graphical user interface (GUI) to install and configure both backup and HSM clients, manage media volumes, monitor library units on the EDM, configure backups, restore data, monitor backup activity, and display reports. Also, use it to set up the migration of data if you have the HSM option.

**Note:** The software for HSM must first be installed directly on the client.

Context sensitive online Help is available throughout the EDM user interface. Online versions of hardcopy books are available through the EDM button in Online Help.

**Note:** The name of the EDM (server) is displayed in the title bar of each EDM window.

For Symmetrix Connect and Symmetrix Path, use the GUI to install clients, configure clients, manage media volumes, and monitor library units on the EDM. You can use the GUI to perform network restores and Symmetrix Path restore operations. You cannot use the GUI for Symmetrix Connect restores.

---

## Starting the EDM GUI

To start the EDM GUI on the EDM server, log in to the EDM. (If you want to configure and manage backups, log in as root or enter `su -` from your user account. The `-` is required.)

Then set your environment and enter the **edm** command:

```
# setenv DISPLAY nodename:0.0
# edm &
```

### Starting the EDM GUI Remotely

From any EDM or client, you can remotely launch the EDM GUI from any other EDM and display it on any EDM or client.

1. Set your environment to designate the machine on which you want to display the EDM GUI.

```
client# setenv DISPLAY nodename:0.0
```

2. Enter **edmremote** with the remote EDM name, and enter the root password when prompted.

```
client# edmremote edm_name
Please enter password for <root> on <edm_name>
Password:
edmremote - Remote GUI launch complete on
<edm_name>. DISPLAY = nodename:0.0
```


Refer to the **edmremote** man page for variations.



## Main Window

The EDM Main window provides a central location for viewing and managing clients and library units connected to the EDM server, and for monitoring and reporting on backup activity.

The main viewing area contains the browser and browser overview, as shown.



You can access the following EDM functions from the menu bar, tool bar, or a pop-up menu in the browser.

- Backup Client Install Wizard
- Backup Configuration Wizard
- Backup Activity Wizard
- Volume Management
- Backup Configuration
- Backup Report
- Restore
- System Monitor Configuration
- HSM Client Installation<sup>1</sup>
- HSM Configuration

**Note:** The HSM Client Installation and Configuration windows do not appear unless you have the HSM option installed on your EDM.

---

## Backup Client Install Wizard



Click on this icon to install a backup client. The Backup Client Install Wizard leads you step by step through the install process:

Refer to EDM online help for more information about this wizard.

When you complete the installation, a window opens and asks if you want to configure a simple backup.

---

## Backup Configuration Wizard



The Backup Configuration Wizard enables you to configure a network, Symmetrix Path, or Symmetrix Connect backup of filesystems or a database. This wizard supports database backups for Oracle, Informix, Exchange, SQL Server, Lotus Notes, and Sybase. The Backup Configuration Wizard leads you step by step through the configuration process. Refer to EDM online help for more information about this wizard.

When you complete the configuration, backups can start, either automatically if you choose that option in the wizard, or manually using either the Backup Activity Wizard or the command line interface.

---

## Backup Activity Wizard



The Backup Activity Wizard enables you to start new, queued, or failed backups, stop running backups, or manage the backup queue.

In the Wizard panels you select a backup operation, select the objects that you want to operate on, choose backup options, and confirm your actions. Then from the Main window, you can monitor the progress of the backup operation that you initiated. (Refer to EDM online help for more information about this wizard.)

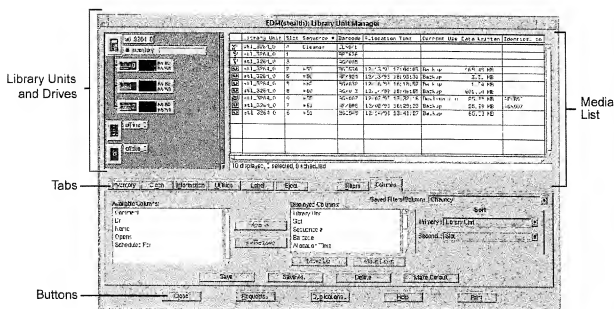
**Note:** You must have root privileges or be an EDM Backup Administrator to use the Backup Activity Wizard.

## Volume Management



Click on this icon to display the Library Unit Manager window. In this window, you manage and monitor media, library units, and drives. You can label media (volumes), take an inventory of tapes or optical disks in a unit, find specific tapes or disks, create and save customized settings in the media list, restart failed media duplications, and perform other related tasks. (Refer to EDM Online Help for more information about these options.)

Volume management also alerts you when operator intervention is needed. For example, if additional media is needed to complete a backup, a window appears that indicates the volume(s) that are needed to complete the operation. For more information about volume management, refer to Chapter 7 "Basic Volume Management Concepts," Chapter 8 "How Volume Management Works," and Chapter 9 "Media Duplication."

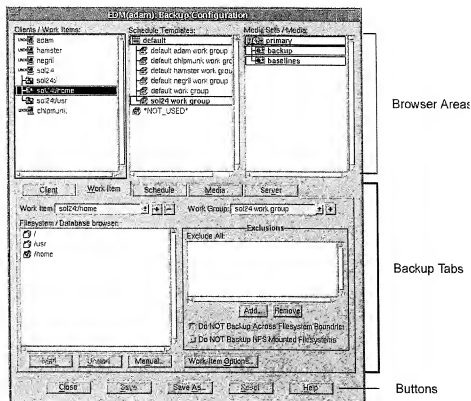


## Backup Configuration



Click on this icon to display the Backup Configuration window. Use it to view backup schedules, configure clients for backups, assign media sets to backups and assign backup administrators to perform restricted tasks.

You can also set backup schedules based on parameters that you specify for your site. This configuration process enables you to tailor the backup system to meet the needs of your organization. For an understanding of backup and restore, review Chapter 3 “Basic Backup and Restore Concepts,” Chapter 5 “How Backup and Restore Work,” and Chapter 6 “Database Backup and Restore.”



## Backup Report



Click on this icon to display the Backup Report window, where you can execute reports on active and completed backups. You can view reports on the local EDM, or participants of an EDM domain can view domain reports on other EDMs in the domain.

This window enables you to create, modify, and save backup reports in several key areas, such as performance within specified time periods, work items with poor performance, or failed work items.

You can configure a report that filters on time and date ranges, work item characteristics, backup states, throughput information, and backup level. You can also designate the columns the report displays.

**EDM Workbench: Backup Report**

File Edit Domain Options Help

Local Reports Available

- Active Backups
- Active and Disabled
- Local Local Backups**
- Local Backups

Report Run Time: Tuesday, August 11, 1998, 12:04

Work Item	Level	Date	Time	Start Date	Run Time	State
WorkItemLocal_00000000	1	08/08/98	08:00:00	08:00:00	08:00:00	Missing or inconsistent configuration file
WorkItemLocal_00000001	1	08/08/98	08:00:00	08:00:00	08:00:00	Missing or inconsistent configuration file
WorkItemLocal_00000002	1	08/08/98	08:00:00	08:00:00	08:00:00	Missing or inconsistent configuration file
WorkItemLocal_00000003	1	08/08/98	08:00:00	08:00:00	08:00:00	Missing or inconsistent configuration file
WorkItemLocal_00000004	1	08/08/98	08:00:00	08:00:00	08:00:00	Missing or inconsistent configuration file
WorkItemLocal_00000005	1	08/08/98	08:00:00	08:00:00	08:00:00	Missing or inconsistent configuration file
WorkItemLocal_00000006	1	08/08/98	08:00:00	08:00:00	08:00:00	Missing or inconsistent configuration file
WorkItemLocal_00000007	1	08/08/98	08:00:00	08:00:00	08:00:00	Missing or inconsistent configuration file
WorkItemLocal_00000008	1	08/08/98	08:00:00	08:00:00	08:00:00	Missing or inconsistent configuration file
WorkItemLocal_00000009	1	08/08/98	08:00:00	08:00:00	08:00:00	Missing or inconsistent configuration file
WorkItemLocal_00000010	1	08/08/98	08:00:00	08:00:00	08:00:00	Missing or inconsistent configuration file
WorkItemLocal_00000011	1	08/08/98	08:00:00	08:00:00	08:00:00	Missing or inconsistent configuration file
WorkItemLocal_00000012	1	08/08/98	08:00:00	08:00:00	08:00:00	Missing or inconsistent configuration file
WorkItemLocal_00000013	1	08/08/98	08:00:00	08:00:00	08:00:00	Missing or inconsistent configuration file
WorkItemLocal_00000014	1	08/08/98	08:00:00	08:00:00	08:00:00	Missing or inconsistent configuration file
WorkItemLocal_00000015	1	08/08/98	08:00:00	08:00:00	08:00:00	Missing or inconsistent configuration file
WorkItemLocal_00000016	1	08/08/98	08:00:00	08:00:00	08:00:00	Missing or inconsistent configuration file
WorkItemLocal_00000017	1	08/08/98	08:00:00	08:00:00	08:00:00	Missing or inconsistent configuration file
WorkItemLocal_00000018	1	08/08/98	08:00:00	08:00:00	08:00:00	Missing or inconsistent configuration file
WorkItemLocal_00000019	1	08/08/98	08:00:00	08:00:00	08:00:00	Missing or inconsistent configuration file
WorkItemLocal_00000020	1	08/08/98	08:00:00	08:00:00	08:00:00	Missing or inconsistent configuration file

Total Work Items: 1,000,000 Total Time: 17:00:00

Report Tabs: Basic Advanced Details Auto-Report

Objects: [List of objects]

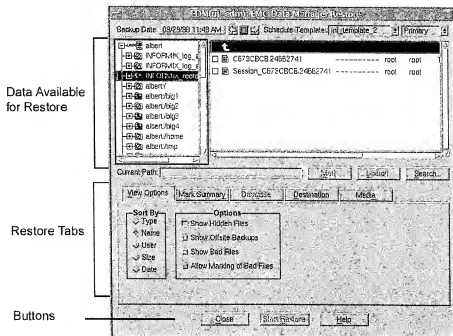
Date: Entire History Date Range From To Status: All Status Succeeded Failed In Progress Occurrences: All Occurrences All Errors All Warnings

## Restore



Click on this icon, or enter **edmrestore** on the server command line, or enter **edmcrestore** on a client command line to display the Restore window. Use it to restore data that the EDM has backed up.


To simplify restoring files, catalogs of backed up files are maintained on disk. Using the restore program, you can browse on-line catalogs, mark files or entire directories for restore, and restore them to a selected client. For an understanding of restore, see "How Restore Works" on page 5-16.



## System Monitor Configuration

In the GUI, you can configure many RASD and RSM parameters. RASD creates various alerts that can be sent by email to a unique list of recipients, which includes Customer Service personnel and/or the system administrator, an SNMP trap or Tivoli Foundation Level event notification to a designated workstation, or calls home to the EMC Customer Service Database system. The type of alert depends on the severity of the problem, the alert recipients, the length of time that the problem persists. Functions that RASD monitors are scalable and configurable. RASD functionality enables monitoring of items such as available volumes using watermarks, cleaning cartridges using watermarks, drives that need cleaning, unsatisfied volume requests, swap space and filesystems using watermarks, failed backups, client availability, patches and firmware revisions, failed duplications, read/write errors, etc.

The System Monitor Configuration GUI interface is part of the EDM Main window. Under the Main window, choose **System Monitor > Configure...**




Also part of the EDM RASD software is the ability of Tivoli Foundation Level integration so that Tivoli customers can call the GUI from their framework package. Administrators can launch the EDM GUI from Tivoli, customize their management GUI to call the EDM GUI, provided that their Tivoli Event Console (TEC) is a valid EDM backup client. You can learn more about RASD through the GUI online help.

## HSM Client Installation



If you have the HSM option, click on this icon to display the HSM Client Installation window to make an HSM client (on which HSM software was previously installed through ep\_install) visible and configurable from the HSM Configuration window.

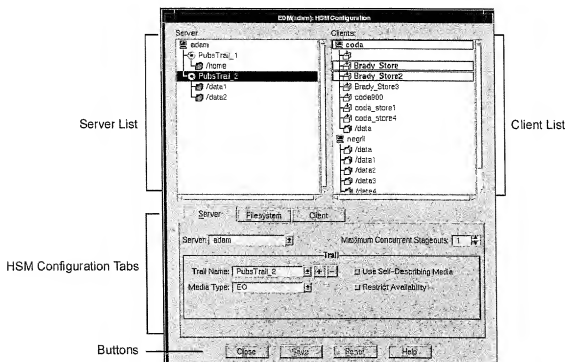




## HSM Configuration



If you have the Hierarchical Storage Management (HSM) option, click on this icon to display the HSM Configuration window to set up stageable filesystems on the server and clients. This window enables you to set up the details of migrating file data out to optical disks or magnetic tapes, creating a much larger virtual filesystem. For detailed information about HSM, review Part II, "Hierarchical Storage Management (HSM)," in this manual.



---

---

## 2 Managing Your EDM System

The *EDM Software Reference* provides you, the system administrator, with step-by-step instructions for managing the EMC Data Manager (EDM) system with or without the Hierarchical Storage Management (HSM) option.

This chapter outlines the procedures that you perform regularly to manage your system. Topics in this chapter include:

- Administering the System
- Running Procedures Automatically via Cron
- Backing Up Server Database Files

---

## Administering the System

With an EMC Data Manager system, you should perform several tasks at the beginning of the day.

Most repetitive tasks are set up to run automatically from root's crontab file. This section briefly describes both manual and automatic tasks.

If you have the HSM option, remember that even though HSM software can migrate files to and from staging media, the files must be backed up, whether they reside on the server's magnetic disks or on the staging media. As with any data storage solution, backups are a must.

As a system administrator, you should perform the following tasks on a regular basis:

- ❑ Check daily to ensure an adequate supply of media is available in the library unit(s) for the next backup and migration runs.
- ❑ Verify that at least one cleaning cartridge is available in the library unit.
- ❑ Check the Backup Report window in the EDM graphical user interface (GUI), or the history report, to ensure that your daily backups (and duplications, if any) completed successfully.
- ❑ Read the daily message log file to review system activity.
- ❑ Run and save the backup disaster report.
- ❑ With the HSM Option, check compaction of staging and baseline volumes (weekly).

Database/system administrator duties for Symmetrix Connect backup and restore activities are described in the *EMC Data Manager Symmetrix Connect User Guide*.

---

## Checking Media

Every day you should make sure that enough media is available in the library units to perform scheduled tasks such as the nightly backup and periodic staging runs. Always keep a supply of pre-labeled volumes in the library unit. Also ensure that a cleaning cartridge resides in the library unit. Refer to Chapter 7 “Basic Volume Management Concepts” for more information.

---

## Verifying Backup Completion

You can check on completed backups at any time by viewing the Backup Report window in the EDM GUI. Access this window from the Main window by clicking on this icon in the Main window tool bar.



In the Backup Report window, you can view and execute reports on the local EDM or on a group of EDMs that are clearly defined as a Domain. The report can include backup attributes for a specific work item such as its backup status, total throughput, total size of the backup, total files that were backed up, and others. For detailed information about the Backup Report window, click the Help button in the window.

You can also run **ebreport backup** every day, after the daily backup is scheduled to complete, to verify that all of the scheduled backups completed.

The following example lists all of the backups that have run since a specified date:

```
# ebreport backup -since date
```

The following example lists all of the backups that have run for a particular client since a specified date:

```
# ebreport backup -client clientname -since date
```

Refer to “Backup Duplicate Reports” on page 16-12 and the **ebreport** man page for more information.

## Reading the Daily Message File

The error logging facility produces a log file on a daily basis that reports system activity for the past 24 hours. This log file is located in `/var/adm/epoch/daily`. (See Chapter 15 “Message Logging” for more information.) If you have a mail facility, you can edit root’s crontab file to mail this log file to the appropriate users.

The portion of the crontab file that describes the daily log file follows:

```
# Daily log
# eptrunclog truncates /var/adm/epoch/daily to a one day slice.
# eptrunclog can also mail the log to one or more specified users.
# eptrunclog should not be run at the same time as epnewlog
# to mail the log to different user(s), replace "root" with the user list
# to skip mailing the log, delete "root" from the command line
00 7 * * * /usr/epoch/lib/eptrunclog "root" > /dev/null 2>&1
```

By default, the `trunc_daily_log` script mails the log file to root. To send the log file to additional users, do the following:

1. Log in as root.
2. Open root’s crontab file with the editor of your choice.
3. Locate the line:

```
00 7 * * * /usr/epoch/lib/eptrunclog "root" > 2>&1
```

4. Add the names of the users to whom you want to send the log files. Separate user names with a space. For example:

```
00 7 * * * /usr/epoch/lib/eptrunclog "root chr" > 2>&1
```

5. Save the file and exit from the editor.

---

## Saving the EDM Backup Disaster Report

At the completion of every backup, the `/usr/epoch/EB/config/local_db_cleanup` script automatically generates a MINIMAL Disaster Report. By default, this report is e-mailed to all EDM Backup administrators, appended to `/usr/epoch/EB/config/disaster-report.log`, and printed to the default system printer.

**CAUTION: It is essential that, for each backup, you save a hard or soft copy of the MINIMAL Disaster Report in a fireproof location, either offsite or in an onsite fireproof vault.**

This MINIMAL Disaster Report is a subset of the FULL Disaster Report generated by **ebreport disaster**. It provides essential information that you need to perform a disaster recovery on the server — a list of media volumes for the most recent LOCAL\_DATABASE backup, the current EDM Backup configuration, the current Library Manager configuration, copies of the key configuration-file settings, and information about baseline backups.

This MINIMAL Disaster Report does *not* include backup client information.

You should run the FULL Disaster Report once every backup rotation and whenever significant system changes were made. The following example runs the FULL Disaster Report and redirects it to a file:

```
# ebreport disaster > ~sysadmin/disreports/960917
```

Refer to Chapter 19 for more information on being prepared for a system disaster. See “Backup Disaster Reports” on page 16-19 for a description of the FULL Disaster Report.

---

## Check Compaction of Staging and Baseline Volumes

In the HSM option, compaction is a collection process that, in effect, frees up staging volumes and baseline volumes in a library unit, which ensures a pool of available media.

You can configure HSM software to compact staging volumes automatically. (See the **emcompact** line in root's crontab file.) If you use baseline backup, you need to compact your baseline volumes manually. Note that only reusable media can be compacted automatically. Refer to "Compaction of Staging Media" on page 11-27 and "Compacting Baseline Media" on page 11-29 for details.

---

## Running Procedures Automatically via Cron

You run most repetitive procedures automatically from root's crontab file (refer to on page 2-7). The **cron** entries are placed in root's crontab file during configuration. If you do not use the autoconfiguration option, the configuration programs prompt you for crontab entries.

**Note:** Adding new entries to the crontab file through the EDM GUI does not replace existing entries with the same characteristics. You must remove existing entries from crontab manually. It is recommended you do this by running **crontab -e** as root. (Refer to the crontab(1) man page for more information.)

You can also schedule a backup in the crontab file within the Backup Configuration window of the EDM GUI. In the Schedule tab, you choose a work group for backup; select Schedule in CRON in the CRON Options section, and enter the time that the backup is to occur. You can also indicate whether you want to retry a failed backup, or use new media for a backup. EDM places this information in the crontab file; the backup then runs automatically at the specified time, on a daily basis.

For maximum efficiency and backup coverage, EDM Backup default settings back up the server first, followed by the network clients, and finally, the EDM Backup and Volume

Management databases. You can back up the server, several clients, and the server database all within a single backup template.

The following entries are added during EDM Backup configuration:

```
# Entries for EpochBackup end with this comment: #EPCebs
# Invoke EpochBackup backup program #EPCebs
00 18 * * * /usr/epoch/EB/bin/ebbackup default >/dev/null 2>&1 #EPCebs
# Invoke EpochBackup catalog cleanup program #EPCebs
00 1 * * * /usr/epoch/EB/bin/ebcatclean >/dev/null 2>&1 #EPCebs
# Invoke EpochBackup catalog expiration program #EPCebs
00 11 * * * /usr/epoch/EB/bin/ebexpire -expire -purge >/dev/null 2>&1
#EPCebs
```

The following table lists procedures that are often run from **cron**. Note that most crontab lines that are used to invoke procedures must explicitly set the full path.

Table 2-1

cron Procedures

Procedure:	Frequency:	For further information:
Run <b>epnewlog</b> to rotate, archive, or truncate EDM system logs.	Hourly and weekly	See "Rotating Error Logs" on page 13-8 and the <b>epnewlog</b> man page.
Run <b>emvck</b> to check and correct staging volume statistics (HSM only).	Daily	Refer to the <b>emvck</b> man page and refer to the daily message log.
Run periodic staging (emmasterd.pid line in root's crontab (HSM only)).	Daily	See "Filesystem Configuration and Maintenance" on page 11-3 for details.
Run <b>emcompact</b> to compact staging volumes automatically (HSM only).	Daily	Refer to "Compaction of Staging Media" on page 11-27 for details.
Run <b>eptrunclog</b> to truncate the daily message log file and optionally, to mail a copy to specified users.	Daily	See the <b>eptrunclog</b> man page.



Table 2-1

cron Procedures (Continued)

Procedure:	Frequency:	For further information:
Run <b>epcleanup</b> to remove files that are no longer needed.	Daily	See the <b>epcleanup</b> man page.
Run <b>ebbackup</b> to back up your server, EDM Backup clients, and EDM Migration clients.	Daily	Refer to "Backup Processing" on page 14-2, and the <b>ebbackup</b> man page.
Run <b>ebreport disaster</b> once every rotation period to keep track of backup media. (A minimal disaster report is generated automatically after each backup.) This report should be stored in a safe place (on another system and on hard copy) for use in recovering from a disaster.	Daily	Refer to "Backup Disaster Reports" on page 16-19 and the <b>ebreport</b> man page.
Run <b>ebreport history</b> after every backup session to see which systems were successfully backed up. This can be included in a short script which also sends mail to the user community.	Daily	Refer to "Backup Duplicate Reports" on page 16-12 and the <b>ebreport</b> man page.
Run <b>ebbackup -drain</b> together with <b>ebbackup -halt</b> if you want to ensure that backup terminates at a certain time.	Daily	See the <b>ebbackup</b> man page.
Run <b>ebcatclean</b> to delete incomplete backup catalogs that may have been created by failed backups.	Monthly	See the <b>ebcatclean</b> man page.
Run <b>ebexpire</b> to manage expiration of catalogs, media, saveset records, and incomplete backups.	Weekly	See the <b>ebexpire</b> man page.
Run <b>ebexpire -purge</b> to delete expired catalogs, media, saveset records, and incomplete backups.	Weekly	See the <b>ebexpire</b> man page.
Run <b>ebexpire -partial</b> to delete incomplete catalogs and backups.	Weekly	See the <b>ebexpire</b> man page.
Run <b>ebexpire -list_orphans</b> to display a list of orphaned volumes.	Varies	See the <b>ebexpire</b> man page.
Run <b>ebexpire -free_orphans</b> to display the list of orphaned volumes and then deallocate them to make them available.	Varies	See the <b>ebexpire</b> man page.

Table 2-1

cron Procedures (Continued)

Procedure:	Frequency:	For further information:
Run <b>emcheck</b> to clear incomplete bit files from client stores (HSM only).	Daily	See the <b>emcheck</b> man page.
Run <b>emsundel</b> to recover bit files from the server's backup volumes. (HSM only).	Daily	See the <b>emsundel</b> man page.
Run <b>evmclean</b> to clean tape drives.	Varies	See the <b>evmclean</b> man page. Also, refer to the tape drive's manufacturer for details regarding maintenance scheduling.

## Deleting Existing Entries in the crontab File

Adding new entries to the crontab file through the EDM GUI does not replace existing entries for the same activities. You must remove existing entries from crontab manually. After creating new, equivalent entries through the GUI, it is recommended that you delete existing entries by running the command **crontab -e** as root (refer to the **crontab(1)** man page for more information).

Note that the EDM GUI controls pre- and post-commands creation; thus, you do not have direct control over those extensions to the ebbackup command when you are using the GUI. The GUI does not allow you to configure complex or non-standard pre- and post- commands. If you want to use complex or non-standard pre- and post- commands, use the **crontab -e** command.

---

## Backing Up Server Database Files

All configuration, backup, and volume information resides in individual server database files. EDM Backup and Volume Management each maintains its own information. The server database files consist of:

- volume management database
- backup catalogs
- backup management files

The LOCAL\_DATABASE work item, which is created as part of server autoconfiguration, includes the pathnames of the server database files. It is essential that this work item is part of the nightly backup. The server database files are essential for performing a complete restore of the server and clients in the event of a disk failure.

Always back up the EDM Backup databases on the server after you back up the clients. This is also the case even if you have no network clients, because the server is considered a "local" client to EDM Backup.

Database backups provide you with complete information about both the server and client backups and shorten disaster recovery time because they allow you to restore the database independently from the files that were already backed up. By default, the LOCAL\_DATABASE work item is backed up last.

If the LOCAL\_DATABASE work item remains in the schedule for more than 24 hours without being run, it is forced to run immediately. This is known as a "late" LOCAL\_DATABASE backup.

For more information about the backup and volume management files which make up the server database, refer to Appendix A "Directory Structure". Backup catalogs are described in "Cataloging of Backup Data" on page 3-11.

---

## 3 Basic Backup and Restore Concepts

The EDM Backup software automatically backs up computers throughout your network. It works with volume management software, which manages storage media in robotic library units.

This chapter describes the basic concepts of backup configuration and operation. Its focus is on filesystem backup over the network.

The topics in this chapter include:

- Client/Server Architecture
- Key Processing Concepts
- Configuration Options
- Reports and Logs
- Manual Operations

For information on database backup, see Chapter 6, “Database Backup and Restore”.

For overview information on EDM Symmetrix Connect backup and restore, refer to the *EMC Data Manager Symmetrix Connect User Guide*, the *Oracle Backup Client* guide, and the *EMC Backupint* guide.

---

## Client/Server Architecture

Backup and restore software has a client/server architecture:

- *client* software runs on the server (the EDM) and on each client that you want to back up in your network
- *server* software runs centrally on the EDM that also runs the volume management software

Server software automatically administers backups of the data on clients throughout your network and of data on the EDM server itself. It does so according to general parameters shipped with the system and added when you first set up the system. You can change these backup parameters to meet the specific needs of your site.

Remote client software, which is located on each client, receives instructions from the server, scans filesystems, and sends the backup data to the server. Local client software, which is located on the server, backs up the server data to the server's tape library unit.

For level-9 incremental backups, the client software determines which files changed since the last full backup and backs up only those files.


If you purchase the optional EDM Online Database Backup software, you can back up a database without shutting it down.

Without the online option, the basic EDM Backup client software enables offline database backup by shutting down databases prior to backup. After database backups are finished, the client software puts the database back online.

During the restore process, the client software receives the restored data and writes it to the client disk.

For a discussion of Client/Server processing, see "Client/Server Processing Methods" on page 5-6.

Figure 3-1



## Key Processing Concepts

EDM Backup scheduling and processing is automatic and dynamic. Workloads are balanced for speedy and efficient backups every night while creating a manageable set of media for every few weeks' worth of backups.

The data you want to back up is specified in *work items*. The data can be a filesystem, disk, or partition on a UNIX, NetWare, OS/2, or Windows NT platform or an Oracle, Sybase, Informix, or SAP R/3 database. (Refer to the *EMC Data Manager Software Release Notes* for a current list of available clients.)

Work items are collected into *work groups*. A *backup schedule template* specifies where to write your backups and how work items are to be scheduled for backup. You can create separate trails for full and incremental backups. Trails are grouped into *trailsets* (media sets), which specify trails for all of the backup levels that the schedule template uses.

This section discusses the following concepts:

- Scheduling
- Nightly Backup Processing
- Restore Processing

---

## Scheduling

EDM Backup offers several ways to schedule backups:

- Automatic Scheduling
- Custom Scheduling
- Command Line Scheduling
- Backup Activity Wizard

If any client is unavailable for backup, EDM Backup continues to back up the other clients in the work group. EDM Backup automatically reschedules failed clients and balances the entire schedule.

---

### Automatic Scheduling

Automatic scheduling of filesystem backups performs some number of full backups each day for the scheduled period. EDM Backup calculates its own schedule for performing full (level 0) and incremental (level 9) backups. With automatic scheduling you can also change the schedule to perform full backups only during weekends. Configure automatic scheduling in the EDM Backup Configuration window.

**Note:** EDM-initiated database backups are automatically scheduled, too. But with database backups, no determination of full or incremental is made by the EDM backup scheduler. The EDM kicks off the database backups on the database client. The backups are performed according to configuration on the client. See "Database Network Backup Overview" on page 6-7 for more information.

---

### Custom Scheduling

Custom scheduling through the EDM Backup Configuration window enables you to perform filesystem backups other than the levels 0 and 9 backups that automatic scheduling provides. With custom scheduling, you explicitly specify the days and levels of backups for individual clients.

---

## Command Line Scheduling

Command line scheduling enables you to enter command overrides to the schedule template configured by the previous methods. You can use command line scheduling to resume an incomplete backup operation or to manually run a backup that is not currently scheduled with the automatic or custom methods.

---

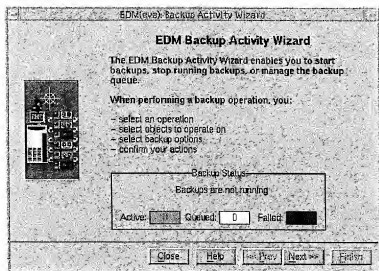
## Backup Activity Wizard



Click this icon in the Main window toolbar of the EDM GUI to access the Backup Activity Wizard. This wizard enables you to start new, queued, or failed backups, stop running backups, or manage the backup queue.

**Note:** You must have root privileges or be an EDM Backup Administrator to use the Backup Activity Wizard.

Refer to EDM online help for more information about the Backup Activity Wizard.






### Concurrent Work Item Input

Backup software enables you to gather backup data from many sources at once. The server can concurrently process multiple streams of data from numerous clients.

For a typical client disk, the entire disk, with all its filesystems, is designated as a single backup *work item*. Each work item is backed up by a single client process, sending a single stream of data to its corresponding server process.

Software also considers various limits on backup processing, such as the total amount of backup data in the network, the total concurrent backup streams allowed to be sent to the server at any one time, and the total concurrent backup streams to be funneled together to the storage media.

Figure 3-2



Large filesystems are designated as separate work items to prevent streams from being too long. (Special parameters prevent conflicting concurrent backups of data from the same disk.)

---

### Balanced Scheduling

Because numerous work items can be scanned for backups separately, an optimal subset of work items can be scheduled for a full backup each night, while the rest of the work items receive incremental backups. Server software intelligently schedules cyclical full backups along with nightly incrementals. This is known as *autoscheduling*.

If you choose to autoconfigure your system, you have the following configuration:

- Every client is backed up according to a single backup schedule template. The backup template lists the default work group.
- Work items are created for each client and inserted into the default work group.
- Backups are scheduled automatically, with each client receiving at least one full backup every two weeks (the rotation period) and receiving an incremental backup on all other nights.
- Backups are written to one media set (trailset) every night.
- All of the data that is sent to the trailset is written to a single media volume or series of volumes (a single trail). (A media volume is a labeled tape cartridge or one side of an erasable optical disk.)

---

### Rotation Period

Software rotates full backups among the work items so that all work items receive a full backup at least once within a *rotation period* (the default is 14 days).

---

### Load Balancing

Every work item is scheduled for either a full or incremental backup each night. To create the nightly backup schedule, the server software considers not only the rotation period, but also actual backup results from previous nights. If all work items received a full backup in the rotation period, the software's *load*

*balancing* feature will schedule a new full backup for some work items, to continue to smooth out the backup work load for each night.

---

## Multiplexed Storage

While numerous work items can be scanned for backups separately, it is possible to multiplex (funnel) the backup data together when writing to the storage media.

## Trails


By default, the server software sends all of the backup data in a single stream of data to a single backup drive, where the data is written to the backup volume. Whenever the volume fills, a new volume is automatically inserted into the drive. The next night, another single stream of data is again written to the volume. At the start of the next rotation period, a new volume is automatically inserted in the drive.

The single media volume or serial set of media volumes written to over the course of one rotation period is called a *trail*.

---

Figure 3-3

Single Trail




**Trailsets**

A complete set of full and incremental backups for a rotation period is called a media set (trailset).

By default, the single trail also constitutes the trailset. But, you could write full and incrementals in two separate streams to two separate trails (sets of media). You would need to use two drives for writing the backups concurrently.

The combined set of full and incremental trails constitutes the trailset. Each trailset contains at least one full backup for each work item (from various nights in the rotation period) plus the incremental backups for every other night in the rotation period.

**Figure 3-4****Multiple Trails in a Trailset**

---

## Nightly Backup Processing

Server software runs automatically every night to administer network-wide backups. The software includes configuration parameters describing what data gets backed up from the magnetic disks of the server and any client computers. The software also maintains status information about what data was or was not backed up recently.

Software refers to configurable *backup shift* guidelines for the maximum number of hours that the entire nightly backup has to start backups. One guideline applies to weekdays, another to weekend shifts.

Client software scans the filesystems and databases (both filesystems and raw partitions) on the hosts and streams the data to the server.

The server writes the data to tape library units attached to it through SCSI connections.

Server software creates online catalogs of the backups to make it easy to restore whenever necessary. Notifications by email inform you of the status of your backups. You can run reports as well.

---

## Automatic Start

A main backup process is started nightly by **cron** from the root crontab file. (For more information, see "Running Procedures Automatically via Cron" on page 2-6.) This process consults the configuration parameters and status information, and then automatically schedules backups for each client's magnetic disks.

The main process spawns separate processes to handle the backup work for each client or portion of the client's data. The individual server processes each send instructions to their corresponding client, and specifies which filesystems and files to back up.

---

## Client Processing

On each client computer (and on the server computer itself) the client software scans the local filesystems as directed.

When prompted for a *full backup* (level 0), the client software copies data for all filesystems, directories, files, and databases specified for backup and streams the data to the server.

When prompted for an *incremental backup* (level 9) the client software sends file data for only those files that have changed since the previous backup. (Unlike the UNIX **dump** level 9, each consecutive level 9 backup copies only files that have changed since the last level 9 backup.)

Whether it is doing a full or incremental backup, the client software always sends complete directory information, so that during the restore process, you can browse an accurate view of the directories as they were at the time of any backup.

---

## Backups of Changing Files

You can continue to work during backups. As the backups are processing, the server software checks the backup directory information to find files that have changed during the backup. Any such files are backed up and checked again two more times that session. (Any files that continue to change both times will be backed up during the following backup session.)

---

## Storage of Backup Data

As the server software receives the backup data, it writes all the data (from multiple clients) to one or more drives in the Library Unit. By default, it streams all of the data to a single backup drive, writing the backup data to the volume in that drive.

---

## Cataloging of Backup Data

Backup software creates *catalogs* that keep track of filenames, file attributes, and locations of backup data. It copies the catalogs onto the media along with the backup data. When a backup completes, the software processes the catalogs and keeps them online so that you can quickly retrieve the data.

To maintain sufficient magnetic disk space for backup catalogs, you'll need to expire the older catalogs after a fixed length of time. The catalog expiration period must be at least one day longer than two rotation periods. To determine catalog expiration, you must consider other system factors and user requirements for executing restores. Catalogs have dependencies on backups and other catalogs. Before you decide on an expiration schedule, refer to "Expiration of Backups and Catalogs" on page 10-2 for more information.

---

## Restore Processing

Anytime you need to restore a file from the backup media, you or your users can use the **edmcrestore** command on the client to open the Restore window in the EDM GUI and display it on the client to browse and restore backed up data. Of course, you can also open the EDM Restore window directly from the EDM Main window on the server.

The Restore window displays file listings derived from the online catalogs. You can browse through directories as they existed on each backup date, and you can browse back and forth among backup dates.

When you have selected the files you want, you start the restore. The backup software locates all data, restores it to the client, and logs the restore activity to log files on the server and client computers.

---

## Configuration Options

As system administrator, you can configure backups for both the server and clients centrally from the EDM Backup server by using the Backup Configuration Wizard and the Backup Configuration window in the EDM GUI. You can restore the backed up files using the EDM Restore window.

You can configure various aspects of your backups to meet your site's needs including your network's computers, library units and drives, data, user profiles, and workshift scheduling demands. Use the EDM Backup Configuration Wizard to set the parameters within the server's backup *configuration file* (eb.cfg). Use the Backup Configuration window to customize those settings, if necessary. For a description of the fields in the eb.cfg file, see Appendix B "EDM Backup Configuration File".

---

## Key Configuration Options

Adjusting the configuration is, for the most part, optional. The configuration file is shipped with standard default values that are ready to run and are generally suitable for most sites.

Your only required configuration task is to install clients. In essence, all you need to do is to specify on the EDM server:

*Which* client computers should be backed up?

With clients installed, you can optionally configure these key aspects of how automated backup processing proceeds:


- *What* data on each client do you want to back up?
- *When* should backups run?
- *Where* is backup data written?



**Which Clients to Back Up?**

Use the EDM Backup Client Install Wizard on the server to specify which client computers to back up. The Wizard leads you through a process where you select the clients from a list and install the software from the server over the network to the clients, as shown in Figure 3-5.

**Note:** For PC clients you must first install the client software directly on the PC system as described in the appropriate user guide. Then use the Install Wizard, so that the EDM server recognizes the client.

**Figure 3-5****Installing Client Software****Installation Options**

The EDM Backup Client Install lets you change various installation defaults before you actually install new clients. This is especially helpful for distinctive client platforms.

Installation options include changing accounts and directories for the client software and communication timeouts.

---

### What Data is Backed Up on Each Client?

The specification of backup work for each client consists of one or more work items. The work item description or *directive* specifies filesystems, directories, files, and databases that are included and excluded from backup. The work item directive is an expanded version of a UNIX **find** statement called a **findxcpio** statement. For more information, see Appendix D "findxcpio Directives".

Use the Backup Configuration Wizard to specify the data you want to back up. Then you can use the Backup Configuration window to edit single work items.

---

### When are Backups Scheduled?

The general parameters for backup scheduling are handled within a *backup schedule template*. The schedule template provides parameters that the autoscheduling function uses. It ties them together with the work items to be backed up and the media set (trailset) to be written to. Use the Backup Configuration Wizard to set these values.

### Rotation Period Trade-offs

The rotation period and the various other scheduling parameters are specified in the schedule template. The rotation period specifies a value that is actually used for two purposes:

- Schedule rotation
- Media rotation

The schedule rotation is the fact that the software must perform a full backup on each filesystem work item covered by the template during this period.

The media rotation is the fact that the software appends each nightly backup to the current volume during this period; at the start of a new rotation period, a new volume is started.

The most common rotation periods are 7 days, 14 days (the default), or 28 days. A key consideration in determining the rotation period is that it is faster to write incremental backups than full backups, and incremental backups use up much less

storage media. But because it is faster to restore files from a full backup than from a series of individual incremental backups, you might want to specify shorter rotation periods for data that changes frequently. If your data changes less frequently, you might want to specify a longer rotation period.

### **Assigned Work Groups**

Work items are grouped into one or more *work groups*, whose backups are coordinated by one or more backup schedule templates.

By default, all server work items are assigned to a single work group named "default server work group." All additional client work items are assigned to a single work group named "default work group," and both work groups are assigned to a single schedule template named "default."

### **Use of Additional Backup Schedule Templates**

All client backup work can be coordinated by one schedule template, but if some client backups need to occur more or less often you can assign some of the work to a separate template and specify different rotation periods for each template. This will write the data for groups of clients to separate media.

---

### **Where is Backup Data Written?**

The primary reason for choosing a certain scheduling option may have more to do with the resulting use of media. For example, multiple work groups, schedule templates, and media sets (trailsets) are useful for organizing a site for accounting purposes; by separating the media to which the data is written, you can charge each group for the media that the group uses.

The schedule template names its corresponding trailset. By default, all of the data is written to a single trailset. Therefore, you can create completely separate sets of backups for work groups by creating separate schedule templates and trailsets for different work groups.

And, as is mentioned earlier, you can create separate trails for full and incremental backups without editing the schedule template; you just edit the trailset to create a second named trail to receive the data from your full backups.

**Note:** Moving backup media offsite before its rotation period ends causes the backup that would use that media to fail. You can avoid a failed backup by using new media for the next backup. You configure the use of new media in the Backup Configuration window of the EDM GUI. Select Advanced Options in the Schedule Tab. In the Schedule Options window that appears, select Use New Media When Current (backup media) Is Offsite.

#### Alternate Trailsets

You can schedule two complete sets of backups on alternate nights and store one set of media offsite. This doubles the number of full backups; one full backup is written to each trailset.

In the EDM Backup Configuration window, you can create two complete sets of full and incremental backups that are written to on alternate nights. To configure two alternate sets of backups, you use a single schedule template, but in that template, you specify an *alternate* trailset in addition to the *primary* one.

The template writes to the primary trailset the first night of the rotation period and to the alternate trailset the next night, and so on. Twice as many full backups are run. In addition, there is more incremental data, because, instead of backing up files that changed in the past day, each incremental backs up files that changed in the past *two* days — since the previous backup *to the same trailset*.

#### Custom Schedules

In the Backup Configuration window of the EDM GUI, you can create separate trails for full explicit incremental backup levels (1-8). You create a custom schedule within a template and edit the trailset to create various named trails for one or more levels from 1-8.

In a schedule template, you can override autoscheduling and custom-schedule backups of particular levels on particular days for certain work groups. This should be done as an exception rather than the rule, as you are bypassing the software's autoscheduling intelligence and its benefits.

But you can do this just for special data so that you can be certain that it gets backed up on a set date, so that you can specify one or more *explicit* incremental backup levels (levels 1-8) as well as levels 0 and 9, and so that you can write that data to a separate trail.

### **Media Duplication**

Media Duplication allows you to create a duplicate set of backup media automatically after each backup session.

After you configure media duplication in the EDM Backup Configuration Wizard, the duplication of a set of backup media occurs automatically after each backup session. This background activity starts after nightly backups complete.

For a complete description, see Chapter 9 "Media Duplication."

---

## **Other Configuration Options**

With automatic backups configured, you can also use the EDM Backup Configuration window to configure optionally the following other aspects of the backup and restore software:

- processing of concurrent backups
- identifying individuals who can configure automatic backups, run backups manually, and run restores
- determining the period of time that backup catalogs and data are kept before expiration
- using new media when current backup media is offsite

---

### How are Backups Processed?

More than one backup process can run concurrently on the server and client and be written to one or more backup drives. When extensively tailoring your configuration, you might decide to tune the preset limits on concurrent processing on the server, clients, and to the backup drives.

You can limit these factors affecting backup and network performance:

- maximum number of work items the server can back up at the same time
- maximum number of work items each client can back up at the same time
- maximum number of work items to concurrently write to all trails for each media type
- maximum number of work items to concurrently write to each particular trail (overridden if maximum for that media type has been reached)
- maximum amount of backup data in the network

---

### Permissions and User Modes

You can configure backup to recognize usernames for accounts on your network as a backup administrator or a various class of restore users. (Keep the usernames to eight bytes or less.)

#### Backup Administrator

You can configure EDM Backup to recognize your username as a backup administrator, so that you can run the EDM Backup Configuration window. This also enables you to run backups (**ebbbackup**) manually from the server under your username.

#### Administrator Restores

As backup administrator, you have permission to run restore with full permissions as system administrator from the server or any client computer. You can use the EDM Restore window to browse the backups of any client and change the destination client for the restore.

### Client User Restore Modes

For each client, you can authorize users to use restore with various levels of permissions.

You can configure these in the EDM Backup Configuration window:

- self-service restore users (who only can restore their own files on a single client).
- cross-client restore users (who can restore their own files over to another specified client).
- root-permission restore users (who can restore any files for a single client).

---

### Catalogs and Backup Savesets

The EDM Backup software creates one *backup saveset* for each work item every time it backs up the work item.

The following comprises each backup saveset:

- *backup data*: a copy of each client's backup data on the storage media.
- *backup catalog*: an online listing of the names and attributes of each directory and file in the work item at the time of the backup and the location of backup data for each file that was backed up. Backup processes this catalog after the backup and uses it when needed to restore data.
- *saveset records*: contain information about an entire backup; for example, its start time and the trails that the backup program uses to write the backup data.

Expiration policies for backup data, online catalogs, and online saveset records are defined for the various backup levels. By default, full backups are kept for one year and incremental backups for three months.

You can configure how long to keep backup data (on the backup media) before expiring it so that the media can be reused.

You can also configure earlier expiration of the online catalogs that reference the backup media; you would do this to maintain magnetic disk space on your server, while keeping backup data longer, just in case it is needed. Catalogs take up a lot of magnetic disk space. See Chapter 10 “Magnetic Disk Concepts” for more information.

The backup software needs the saveset records for as long as you keep the backup data. Saveset records are small relative to the catalogs. With the saveset records online, you can use the **ebimport** command if it is necessary to recreate an online catalog for a backup.

---

## Reports and Logs

Backup software maintains logs of all activity, mails notifications about backup processing, and provides various reports to aid you in monitoring the status of backups.

Backup software maintains logs for backups and restores on both the backup server and on each client. It also maintains, on the server, a log that details the backup activities and cataloging operations for each backup template.

EDM Backup automatically sends email notifications about backups that are in progress, that have completed, or failed. You can modify the management of mail notifications and log files in the Backup Configuration window of the EDM GUI.

Refer to Chapter 15 “Message Logging” and Chapter 16 “Backup Reports and Log Files” for details.



---

## Reporting in the EDM GUI

You can monitor active backup processes and execute reports in the EDM GUI.

During a backup, an object in the Main window such as the EDM server, a client, or a work item appears as an active process, successfully backed up, in the backup queue, or failed to complete successfully. Current backup throughput also appears for a backup in progress.

Upon completion of a backup, you can then configure, save, and print backup reports on specific areas of importance such as failed work items or work items with poor performance.



Click on this icon in the Main window toolbar to access the backup report module.

For more information about active backup reporting, refer to EDM online help, "Backup Report Overview."

---

## Manual Operations

Much of the backup operation runs automatically, but you can use certain manual operational and reporting commands at the command line or set them to run automatically from the root crontab file. Refer to Chapter 18 for a list of man pages that are available for backup and restore.

---

## **4** Port Control

EDM port control allows the EDM to communicate with clients on the other side of a firewall. Port control is available for use with UNIX and NT filesystem backups and with UNIX database backups.

This chapter contains the following sections:

- Understanding Port Control
- What is a Firewall?
- Port Control Checklist
- Setting Up the Firewall for Port Control
- Enabling Port Control on the EDM
- Installing Port Control on the EDM Client(s)
- Making Changes to Port Control

Port control allows you to control the TCP ports used by the EDM to communicate with the clients. It also makes network analysis and auditing of EDM network activity easier. It also allows you to take advantage of the router's ability to prioritize packets.

The discussion in this chapter is limited to how port control allows EDM TCP port usage to behave in a predictable manner so that firewall rules can be implemented. A firewall restricts access between networks based on rules set by the local firewall administrator.

---

## Overview

EDM port control allows EDM TCP port usage to behave in a predictable manner so that firewall rules may be written to allow the EDM to communicate with EDM clients on the other side of the firewall.

The port control feature is available for use with UNIX and NT filesystem backups and with UNIX database backups. It gives you the option to control the TCP ports used by the EDM. Port control functionality allows you to control TCP port usage so that you can:

- Back up and restore UNIX and NT filesystems and UNIX databases in a firewall environment
- Analyze and audit of EDM network activity
- Take advantage of the router's ability to prioritize packets across the network, based on your own requirements

EDM port control addresses only a portion of a complete security solution. When properly configured, it eliminates EDM's dependency upon the backup clients' portmapper. Port control must be coordinated between an EDM and all clients that are expected to use it.

---

## Understanding Port Control

By default, port control is disabled on the EDM. To enable port control on the EDM server use **eb\_server\_config** without the **-D** option. You can configure an EDM to have some clients using port control and some clients not using port control. To enable port control for selected client(s) use the Backup Configuration Wizard in the GUI, or the **eb\_client\_install** command with the **-portcontrol** option.

Enabling port control is an easy procedure when implemented correctly at the time of server installation or update. Discuss your firewall policies with the EMC service personnel who install your EDM. To adjust the settings after installation, you must use the **portservices** CLI to change port values, rerun **eb\_server.config**, and reinstall all participating clients (see the **portservices** man page.). Contact Customer Service for assistance.

---

## Installing and Updating Client(s)

Installing, updating, and operating an EDM port control enabled client through a firewall has the same network requirements as other EDM clients, such as network name resolution. Port control allows normal operations to take place in definable tcp port ranges.

Installing and updating a UNIX EDM client requires either UNIX rsh or UNIX rexec accessibility from the EDM to the client and the ability to ping the EDM from the client. These protocols are not usually permitted by firewalls and will need to be allowed during the installation or update.

**Restrictions**

Please note the following restrictions (see Table 4-1 on page 4-6 for default port definitions):

- Port control is not a backward-compatible feature, therefore it requires EDM clients to be updated to EDM 4.5 versions. It supports most versions of UNIX and NT clients (but not Pyramid, Sequent, NCR, SCO and Alpha NT).
- Network database backups require a return connection through the firewall from the client.
- In client-initiated backups, ports have to be open from the client to the EDM.
- EMC does not recommend by-passing firewall policies by bridging the DMZ and trusted network with the Symmetrix. Therefore, if you want to use port control with EDM Symmetrix Path or EDM Symmetrix Connect, you should discuss this carefully with EDM customer service and the local firewall administrator before implementing.

---

## What is a Firewall?


A firewall is a combination of hardware and software applications used to create a gateway and provide controlled access from one network to another. By reviewing all traffic and selectively allowing or disallowing data to pass, the firewall protects the internal network from unwanted external intrusions.

A firewall enforces network policy regarding the restriction of access based upon rules set by the local firewall administrator.

---

## EDM Firewall Assumptions

EDM port control assumes that the Perimeter Network Zone is separated from the EDM by an internal firewall and that it is reasonably secure. If the EDM is in the trusted zone, port control allows it to work within the firewall rules such that the EDM can perform a backup of a client in the Perimeter Zone. Port control can also be set up so that the clients in the perimeter zone can initiate backups.



For normal backup and restore operations, the firewall between the trusted zone and the perimeter zone must be opened to allow TCP communication. The scope of the opening depends on your company policies.

The following are baseline assumptions needed for using EDM port control to backup and restore UNIX and NT filesystems and UNIX databases through a firewall:

EDM port control:

- Assumes an IP filtering-capable firewall (which most firewalls are).
- Uses the TCP protocol. (Use of UDP is not required and ICMP/SNMP are used by RASD for non critical functions.)
- Exists so that rules can be written for a firewall that allow the EDM to interact with EDM clients via TCP. Include the local firewall administrator in the process of defining ports.
- Allows rules to be created within routers to prioritize packets.

**Table 4-1**

**Default Port and Port Range Definitions**

	<b>From Server to Client</b>	<b>From Client to Server</b>
UNIX Filesystem Backup and Restore	8000:8250	—
UNIX Database Backup and Restore	8000:8250	5600
UNIX Client Initiated Filesystem Restore	8000:8250	8000:8250
with GUI Display	X11 ports (6000:6063)	—
UNIX Client Initiated Database Backup and Restore	8000:8250	5600
Windows NT Filesystem Backup and Restore	3895	—

- If you are using UNIX portmapper, you need to open tcp port 111 along with 8000:8250.

- Installing and updating EDM clients through a firewall requires the firewall to be open temporarily. These ports may then be closed after installation/update of EDM client(s).

Some firewalls understand these protocols and allow them to be generically specified to take care of the details.

For those that do not, see Table 4-2:

**Table 4-2**

**TCP Connections for Client Installation and Update**

<b>To Allow:</b>	<b>From The EDM to the Client</b>	<b>From the Client to the EDM</b>
UNIX RSH from the EDM to the client	514	TCP connections on privileged ports less than 1024
UNIX REXEC from the EDM to the client	512	TCP connections on non-privileged ports greater than 1024, but usually greater than 5K, up to 65K
ping from the client to the EDM		ICMP

Note that the UNIX REXEC protocol passes the password, in this case the root password for the client, unencrypted. For this reason, the local firewall administrator might choose to allow the UNIX RSH protocol and temporarily place a `.rhosts` file in the root directory on the UNIX client.



---

## Firewall Requirements

Firewalls control connections. Once connections are established, data can flow bi-directionally. IP firewalls need to know about protocol, source address/port and destination address/port. The side of the firewall the connection request originates from (DMZ or Internal Network) impacts required ports.

Firewalls execute rules in order, looking for the first match. Some typical rules are:

- Allow tcp source <edm> any destination <client> 8000:8250
- Allow/drop/reject matching request.
  - allow - like router, forward packets.
  - drop - silently drop packets (*timeout*).
  - reject - notify originator (*connection refused*).
- Protocols are usually TCP, UDP, ICMP.
- Source (from) what address range/port range.
- Destination (to) what address range/port range.
- Source and Destination are specified as universal addresses:  
<system name or ip>:<port range>

### Examples:

```
myedm.customer.com:8000-8250
193.45.5.25:8000-8250
193.45.5.0:8000-8250
193.45.5.24:6000
```

---

**Sample Firewall Configurations**

The following examples are provided for the local firewall administrator. These examples assume that the EDM is 123.456.78.155 and the EDM client is 123.456.78.170.

**Basic Port Range Example**

To accept TCP connections from the EDM to a client within the defined port range:

```
allow tcp source 123.456.78.155 any destination 123.456.78.170 8000:8250
```

This allows any TCP port on the EDM (123.456.78.155) to connect to TCP ports 8000 through 8250 on the client (123.456.78.170).

**NT Client Example**

To back up a Windows NT filesystem client, you must allow port 3895 on the NT client to accept TCP connections from the EDM:

```
allow tcp source 123.456.78.155 any destination 123.456.78.170 3895
```

This allows any TCP port on the EDM (123.456.78.155) to connect to TCP port 3895 on the client (123.456.78.170).

**Database Example**

To back up a UNIX database client, you must allow port 5600 on the EDM to accept TCP connections from the client:

```
allow tcp source 123.456.78.170 any destination 123.456.78.155 5600
```

This allows any TCP port on the client (123.456.78.170) to connect to port 5600 on the EDM (123.456.78.155).

---

## Port Control Checklist

Before you begin to enable and configure port control, make the decisions in the Port Control Checklist. The local firewall administrator should participate in this process. These decisions will be used to construct firewall rules and to configure the EDM and participating clients.

It is important to do it correctly the first time.

To change it later will require the use of the **portservices** command line to make changes to the EDM server and then add the changes to the client(s) in order to keep them synchronized. (See "Making Changes to Port Control" on page 4-20.)

---

**Table 4-3****Port Control Checklist**

Decisions to Make	Record Decision and Needed Action
<input type="checkbox"/> Name of EDM for port control. (Must be at least EDM 4.5.0)	
<input type="checkbox"/> Decide which EDM clients and/or subnets will be accessed through the firewall. (Must be versions released with EDM 4.5.0 or greater.)	
<input type="checkbox"/> Decide if the default port range 8000:8250 is appropriate or if another port range is preferred.	
<input type="checkbox"/> Decide if you want to use client-initiated backups for any client. If you do, you must be prepared to open the port range from the EDM client in the DMZ to the EDM server. See Table 4-1 on page 4-6 for details.	
<input type="checkbox"/> Note that a low TCP session timeout value could result in failure of mover-aware backups and possibly restores. Either increase the timeout value or only do non-mover backups through the firewall.	

Table 4-3

Port Control Checklist (Continued)

Decisions to Make	Record Decision and Needed Action
<p><input type="checkbox"/> Decide if you will be using the UNIX portmapper on the client (the default) or the EDM portservices file.</p> <ul style="list-style-type: none"> <li>- Using the UNIX portmapper is conservative, but less secure. This is the default, and requires that TCP port 111 must be open through the firewall.</li> <li>- We recommend using the EDM portservices file. This is more secure since you no longer open port 111, but get the port numbers from local files on the server and the client. (There is a risk that services might not be able to be contacted if the configuration becomes unsynchronized by making changes on only one side).</li> </ul>	
<p><input type="checkbox"/> Decide if UNIX database backups will be performed. If so, the firewall must allow port 5600 to be open from the client in the DMZ to the EDM in the trusted zone. See Table 4-1 on page 4-6 for details.</p>	
<p><input type="checkbox"/> Decide if NT filesystem backups will be performed. If so, port 3895 must be open from the EDM to the NT client in addition to the port range. See Table 4-1 on page 4-6 for details.</p>	
<p><input type="checkbox"/> Decide how you want to install UNIX client(s). Select either EDM Transfer Protocol or Remote Shell. Determine TCP port settings. See "TCP Connections for Client Installation and Update" on page 4-7, for more information.</p>	
<p><input type="checkbox"/> Decide if Symmetrix Path or Symmetrix Connect backups will be performed.</p> <p>EMC does not recommend by-passing firewall policies by bridging the DMZ and trusted network with the Symmetrix. Therefore, if you want to use port control with EDM Symmetrix Path or EDM Symmetrix Connect, you should discuss this carefully with EDM customer service and the local firewall administrator before implementing.</p>	

Some optional features may not work if UDP and ICMP protocols are not allowed on the firewall, such as RASD client pings and SNMP alerts.

**Note:** There are 220 reusable ports per client. Database backup work items require 3 ports each plus 1 for each stream. Filesystem backup work items on the client use 5 ports per work item.

After you complete this checklist, proceed with the following:

- Setting Up the Firewall for Port Control
- Enabling Port Control on the EDM
- Installing Port Control on the EDM Client(s)

---

## Setting Up the Firewall for Port Control

After you understand the issues involved in using port control with your EDM, and have completed the “Port Control Checklist” on page 4-10, have the local firewall administrator make all of the firewall adjustments before you configure the server to enable port control.

---

## Enabling Port Control on the EDM

You must configure the EDM server before you install and configure EDM clients. At that time, you will push the definition files out to the client(s) to enable port control.

While it is possible to make some changes later, it is much easier to activate port control when you configure the EDM for the first time. See “Making Changes to Port Control” on page 4-20 for several examples of changes. Therefore, be sure to complete the “Port Control Checklist” on page 4-10 before configuring the EDM.

To configure the EDM, run **eb\_server\_config** without the **-D** option. You must answer **yes** to activate port control. You only need to do this once. Port control will remain enabled when you run **eb\_server\_config** again.

After logging in as root on the EDM server, enter:

```
# eb_server_config
.
.
.
Do you wish to enable port control on the server? <y/n>[ n ] : y

Port Control Information:

Low Port:           8000
High Port:          8250
Lookup Method:      "portmapper"

Enter the low port number for port control:

    (or just press return to accept the default value in square brackets)

[ 8000 ] :

Enter the high port number for port control:

    (or just press return to accept the default value in square brackets)

[ 8250 ] :

You have the following choices for Lookup Methods:

    1) Portmapper
    2) EDM port services file

    (or just press return to accept the default value in square brackets)

[ 1 ] : 2

Port Control files successfully installed on the EDM server.
```

At this point, port control is enabled on the EDM server and you can enable port control on EDM client(s) when you install them.

## Portservices Files

Portservices files are created in the format `edm_services.xxxx` by the **portservices** command in the server's `/usr/epoch/etc/csc` directory. The `csc` directory remains empty until port control is enabled. (See the `portservices` man page for details.)

The files specify which ports an EDM server uses to communicate with its clients and other EDM servers. The presence of these files indicates that port control is configured.


The following examples assume that:

- `edm1` is the EDM server and the local client.
- `client1` and `client2` are EDM clients.
- `.client1_template` contains uncommitted changes for `client1`.

```
-rw-r--r-- 1 root other      816 Jan  7 16:40 edm_services.client1
-rw-rw-rw- 1 root root      824 Jan  9 16:39 edm_services.client1_template
-rw-rw-rw- 1 root other      812 Dec 20 16:20 edm_services.default
-rw-rw-rw- 1 root other      816 Jan  9 16:40 edm_services.localhost
-rw-r--r-- 1 root other      812 Jan 10 09:42 edm_services.client2
lrwxrwxrwx 1 root other    22 Dec 20 16:34 edm_services.edm1 -> edm_services.localhost
```

Note the file extensions in the examples.


1. The .default file is cloned to create the .client1 and the .client2 files.
2. These are then pushed out to client1 and client 2 to enable port control. They are stored on client1 and client2 as the .localhost file and on edm1 as the .client1 and client2 files.
3. The .localhost file on edm1 is pushed to client1 and client2 and named the .edm1 file.





To change the configuration of client1:

1. Use the **portservices** command (see the `portservices` man page) to create an `edm_services.client1_template` file with uncommitted changes on the server.
2. Reinstall client1, selecting yes in the port control window of the Backup Client Install Wizard.
  - a. This copies the `.client1_template` file to client1 and overwrites the `.localhost` file.
  - b. If that is successful, it will then rename the `.client1_template` file to `.client1` on the server and overwrite the existing file.



**Note:** If there is a `.client_template` file, it contains uncommitted changes.

## .localhost File

On both the server and the client, the .localhost file contains settings for the local system.

To view these settings, on either the server or the client, enter:

```
# portservices -disp localhost
```

The following output appears:

```
Port Configuration from file <edm_services.localhost>
```

```
-----
Low Port Number: 8000
```

```
High Port Number: 8250
```

```
Transient Offset: 31
```

```
Fixed Ports      : 8000:8030 (31)
```

```
Reusable Ports   : 8031:8250 (220)
```

```
Socket Option(s): SO_REUSEADDR|SO_LINGER_10
```

```
Lookup Method    : EDM port services file
```

Service	Type	Ident	Offset	Comment
emrpcd	RPC	390000	0	staging daemon (HSM)
vmdaemon	RPC	390001	1	vmdaemon
emsd_1	RPC	390003	3	migration daemon (HSM)
emsd_2	RPC	390004	4	migration daemon (HSM)
ebfsd	RPC	390007	7	ebfs daemon
07dbapicl	RPC	390008	8	EB Database API Daemon
hsmcd	RPC	390009	9	HSM API Daemon
epccommd	RPC	390010	10	File Browser Daemon
edmlinkd	RPC	390011	11	EDM-Link Daemon
bamd	RPC	390012	12	Backup Activity Monitor
edmdispd	RPC	390015	15	Dispatch Daemon (EDM Restore)
07dbapi	RPC	390016	16	EB Database API Daemon for Sybase
tfsmd	Socket	390020	20	TFSM Calypso Daemon

---

## Installing Port Control on the EDM Client(s)

After port control has been enabled on the EDM, select Install Client from the Backup menu of the EDM GUI. The Backup Client Install Wizard appears. Select the client(s) you want to install.

---

### NOS Client Access Window

For Windows NT filesystem client(s), enter the username and password in the NOS Client Access window.

---

### UNIX Client Install Method Window

For UNIX clients, proceed to the UNIX Client Install Method window. Follow the decisions you made in the Port Control Checklist.

Be sure that the local firewall administrator has implemented the firewall rules decided on in the Port Control Checklist. These ports may then be closed after the installation/update of the EDM client(s).

**Note:** The EDM Transfer Protocol, which requires a password, uses the UNIX REXEC protocol during UNIX client installation. The remote shell uses UNIX RSH during UNIX client installations. Both methods of installation use EDM Transfer Protocol during normal operations.

**Note:** If the port control configuration for the EDM local host file changes, all port-controlled clients must be updated to get the changes. (See "Making Changes to Port Control" on page 4-20.)

---

### Port Control Window

When you reach the Port Control window, click:

- **Yes** to enable port control if it is not already enabled on the client. If port control is already enabled on the client this will overwrite the settings with the current settings on the EDM server.

- **No** to leave the port control settings on the client as they are. This will not enable, disable, nor update port control on the client.

**Note:** If you do not reach the Port Control window, the server does not have port control enabled. Enable the server using **eb\_server\_config** before proceeding.

Once a client has port control enabled and the configuration saved, it can be reinstalled without specifying port control, unless port control is removed from the client using the **portservices** command (see the **portservices** man page).

**Note:** You cannot install a client using an IP address, then move it behind the firewall, and give it a new IP address.

---

## Making Changes to Port Control

If your port control configuration changes on the server, EDM clients should be reinstalled to specify the new settings. If changes have been made, the `edm_services.client_template` file appears on the EDM.

---

### To Change System Monitoring for Ping Errors

Customers accessing clients through a firewall which does not allow ICMP packets to pass through in both directions may get RASD errors. While the default configuration for RASD is to check the availability of all clients, this setting can be modified through the System Monitor window in the GUI.

---

### To Change the Default Port Range

If you want to change the default range, begin with the EDM, then reinstall the client(s).

1. To change the range from the default of 8000:9250 to 9000:9250, shut down the server, remove the old port configuration, make the changes, activate the changes on the server, and restart as shown below:

```
edm# edmproc -shutdown
edm# portservices -portconf default -low 9000 -high 9250
edm# portservices -activate
edm# edmproc -restart
```

2. Install the `edm_services` files on every port-controlled EDM client(s) as follows:

```
edm# portservices -portconf <client> -low 9000 -high 9250
edm# portservices -copyto <client>
```

---

### An Alternative Method

Open the firewall for REXEC or RSH and reinstall the EDM client(s), selecting "Yes" in the Port Control window. This will change the ranges on the clients to match the new range on the EDM.

Close the firewall for REXEC or RSH.

---

**To Enable Port Control with eb\_server\_config:**

1. To determine if an EDM has port control enabled, enter **portservices -disp localhost**. If there are no edm\_services files, the EDM does not have port control enabled.

2. If you want to enable it, do the following:

```
edm# edmproc -shutdown
edm# load_portfile [options]
edm# portservices -activate
edm# edmproc -startup
edm# rpcinfo -p | grep 3900
```

3. **rpcinfo** displays the following:

390007	1	tcp	39552
390011	4	tcp	8011
390010	2	tcp	8010
390008	1	tcp	8008
390012	2	tcp	8012
390015	1	tcp	8015

4. Compare the right column for 390011, 390010, 390008, 390012, and 390015 with the settings shown by **portservices -disp localhost** (see “localhost File” on page 4-17). For instance, 390011 has a low port number of 8000 and an offset of 11. Add these to get 8011 and compare to the output of **rpcinfo**.

---

### To Turn Off Portmapper on the EDM

```
edm# edmproc -shutdown
edm# portservices -portconf localhost -lookup services
edm# portservices -activate
edm# edmproc -restart

edm# portservices -copyto <every port controlled client>
```

To turn off portmapper on the EDM:

1. On the edm:

2. Install the changed edm\_services files on all of the port controlled EDM client(s).

---

### An Alternative Method

Open the firewall for REXEC or RSH and reinstall the EDM client(s), selecting “Yes” in the Port Control window. This will change the ranges on the clients to match the new range on the EDM.

Close the firewall for REXEC or RCMD.

**Note:** To turn portmapper back on, use **-lookup portmapper** instead of **-lookup services**

---

### To Turn Off Portmapper on Client(s)

```
edm# portservices -portconf <client> -lookup services
edm# portservices -copyto <client>
```

To turn off portmapper on client(s):

On the edm, for each affected client:

**Note:** To turn portmapper back on, use **-lookup portmapper** instead of **-lookup services**

---

### To Set Portmapper Off for New Clients

```
edm# portservices -portconf default -lookup services
```

To set the default for any new port control client, enter the following on the edm:

**Note:** To turn portmapper back on, use **-lookup portmapper** instead of **-lookup services**

---

### To Disable Port Control for a Single Client

To disable port control on a single client, do the following on the EDM:

```
edm# portservices -removefrom <clientname>
```

This removes the edm\_services files from the client and saves the edm\_services file for the client in a .template file on the EDM. If port control is enabled for the client at a later date, the .template file is used, thereby restoring the port control settings to what they were before port control was disabled.

---

### To Disable Port Control on the EDM and All of Its Clients

If you want to disable port control on the EDM and all of its clients, do the following on the EDM:

```
edm# edmproc - shutdown
```

```
edm# portservices -removeall
```

```
edm# edmproc -startup
```

If any client(s) cannot be reached, the edm\_services files on the client are not removed and must be removed manually. The edm\_services files are located in the sub-directory etc/csc under the directory in which the EDM client software was installed (usually /usr/epoch).





---

## 5 How Backup and Restore Work

This chapter provides an in-depth description of what actually happens during the backup and restore processes. The following topics are discussed:

- How Backup Works
- How Restore Works
- Media Management

For information on how EDM Symmetrix Path backup and restore works, refer to the *EMC Data Manager Symmetrix Path User Guide*.

For information on how to perform EDM Symmetrix Connect backups and restores — specifically with the EDM Oracle Application Interface and Filesystem Application (for UNIX clients) — refer to the *EMC Data Manager Symmetrix Connect User Guide*.

For information on how to perform EDM Symmetrix Connect backups and restores with RMAN Proxy Copy (on UNIX clients), refer to the *EMC Data Manager Oracle Backup Client* guide. Similarly, see the *EMC Data Manager EMC Backint* guide for information on using Symmetrix Connect with the SAP R/3 System's SAP Tools (on both UNIX and Windows NT clients).

For information on how to perform EDM Symmetrix Connect backups and restores with the EDM-specific interfaces to NT filesystems and databases:

- *EMC Data Manager Windows NT Backup Client*
- *EMC Data Manager Windows NT Oracle Backup Client*
- *EMC Data Manager Windows NT SQL Server Backup Client*
- *EMC Data Manager Windows NT Exchange Backup Client*

---

## How Backup Works


Use the EDM Backup Configuration window or command line interface to configure your filesystem and online database backups. The configuration process involves identifying and scheduling each client's data for backup. Once configured, backups occur automatically, copying the files from the clients' magnetic disks to the server's storage devices.

When client software is installed on the server, the magnetic disk(s) on the server are handled the same way as the local disks on client file servers and workstations. When the server is backing up its own local data, it is referred to as the local backup client.

The following is an overview of the backup process:

---

Figure 5-1



1. The **ebbackup** program starts one backup daemon (**ebbackupd**) for each scheduled resource (work item).

The **ebbackup** utility is started from root's crontab file at a specified time.

2. The backup daemon connects with network clients via the client's connection method.

The backup daemon then tells the client what data to back up on that client and at what level (for example, full or incremental).

3. Remote and local clients call the **startfind** utility which starts a **findxpio** process and then does a file scan.

The client scans its filesystems and sends the information (file attributes and the files to be backed up) back to the server.

4. The backup daemon on the EDM Backup server receives the information from the clients, stores it on backup media, creates a saveset record for the backup, and puts file name and attribute information in the backup catalog database.

5. As each work item is backed up, the catalog daemon, **ebcatalogd**, processes the backup catalog files for future use by the restore program.

6. The **ebbackupd** daemon updates several server databases, such as the volume management and saveset databases.

For HSM systems, **ebbackupd** updates the database so that it can determine the backup status of each client. The daemon updates the `ntag.list` to reflect correctly the mapping between work item names and EDM Backup/HSM tags; and, for baseline backups, it updates the `saveset-to-baseline` relations database to record the volume(s) and saveset ID of the backup. (For an understanding of HSM, read Part II "Hierarchical Storage Management.")

7. EDM Backup creates backup completion reports to inform you of successful and failed backups.

---

## Monitoring Active Backups

You can monitor active backup processes and execute reports or queries in the EDM GUI. During a backup, an object in the Main window, such as the EDM server, a client, or a work item, appears as an active process, successfully backed up, in the backup queue, or failed to complete successfully. Current backup throughput also appears for an backup in progress.

Upon completion of a backup, you can then configure and save backup queries to report on specific areas of importance such as failed work items or work items with poor performance.

For more information about active backup monitoring and reporting, refer to EDM online help, "Backup Report Overview."

---

## Start of Overall Backup

The **ebbbackup** command, which is run out of root's crontab file, starts one **ebbbackupd** daemon for every scheduled work item. The **ebbbackupd** daemon performs the backup of a work item listed in the work group list of the backup template.

Backups of work items for multiple clients on the network can proceed concurrently. You can identify work items that back up the same physical disks so that they do not run at the same time. This prevents disk thrashing, thus improving the time to complete the backup. You can also assign a priority to each work item to control which ones are to be backed up first.

---

## Client Access

The **ebbbackupd** daemon connects with network clients by using the client's connection method. Remote UNIX and PC clients use the EDM transfer protocol or the remote shell (**rsh**) utility; online database clients use RPC. The daemon accesses the local client (the server) directly. On UNIX clients, it invokes the **startfind** script on the client system.

---

### Work Item Specification

The daemon provides the client with a specification of the filesystems, directories, and/or files to be backed up, as defined in the work item configuration. When doing this, the **ebbackupd** command takes into account the priority at which each work item is processed; the level of completeness of the backup of an HSM system; and any exclusion tags (do not process work item A at the same time as work item B, etc.). For incremental backups, the daemon also takes into account the date/time of the last backup.

---

### Automatic Scheduling

If a client that is listed in the backup schedule template's work group is unavailable, EDM Backup's autoscheduling function continues to back up all of the other clients that are listed in the work group. On the next day that the client is available, EDM Backup performs the backup.

---

### Backup Activity Wizard

The Backup Activity Wizard enables you to start new, queued or failed backups, stop running backups, or manage the backup queue. You access this wizard from the Main window of the EDM GUI.

**Note:** You must have root privileges of be an EDM Backup Administrator to use the Backup Activity Wizard.

For more information about this wizard, refer to EDM online help.

---

## Client/Server Processing Methods

EDM Backup consists of software that backs up the server and various network clients.

- EDM Backup's *server software* manages all networked client backups.

It provides central configuration and administration of your client backups. The server software also enables you or your users to restore files from backup media easily.

- EDM Backup's *client software* runs both on the server system (as a local client) and on each networked client.

When prompted by the server, the client scans its filesystems and sends the server the files to be backed up.

Refer to the *EMC Data Manager Software Release Notes* for a current list of clients.

---

## Standard Client Processing

The client passes to the server the name and attributes for all scanned files, along with the data for those files that are selected for backup. The client sends the header information even for files that are not being backed up, to be able to reproduce the state of the filesystem at the time of the backup.

With the standard backup processing model, the client sends this information using the "standard output" channel of the connection, using an extended **cpio** format (provided as part of the EDM Backup software).

---

## High-Speed Client Processing

For high-speed client processing, multiple data streams are generated. The header information, which contains the attributes, is sent in one stream to the server, which in turn writes the header data to a backup catalog. Another stream that contains the file data is sent to the backup media.

---

## Server Processing

The server writes the header data to a backup catalog and sends the file data to the second server process, which in turn writes the data to the backup media.

---

## Filesystem Backup

Filesystem backups are performed while you are online. Files are monitored as the backups are processed, so if a file changes while it is in the process of being backed up, the backup of that file is rejected and another is scheduled.

To reduce the backup workload, filesystem backups include *incremental* backup as well as full backups. A level-9 incremental backup backs up only those files that changed since their last backup. Each night, by default, the server software schedules full backups for some hosts and incrementals for the remainder. The scheduler rotates the full backups among all of the hosts over a rotation period, which by default is two weeks.

You can restore individual files from filesystem backups.

Baseline backups, which are available with the HSM option, back up your most stable files. From that point on, you perform backups relative to the baseline.

Refer to the *EMC Data Manager Symmetrix Connect User Guide* for information on backing up UNIX, Oracle, and Windows NT filesystems using EDM Symmetrix Connect.

---

## Client Scans Filesystem

**startfind** and **findxcpio** are utilities that are part of the EDM Backup software. The **startfind** client script runs a filesystem scan utility called **findxcpio** to collect backup data from the client. This utility operates through the filesystem interface so it can work while the client's filesystems are active.



EDM Backup can detect that a file is being changed while it is being backed up. If **findxcpio** detects that a file changed, it backs up the file again. The **findxcpio** utility tries to copy the file up to three times before it skips it. Under these conditions, the file is backed up on the next scheduled backup.

For more information refer to Appendix D “findxcpio Directives”.

---

## Database Backup

There are several ways to back up databases, as described in Chapter 6, “Database Backup and Restore” of this manual.

**Note:** For Symmetrix Connect, see also the *EMC Data Manager Symmetrix Connect User Guide*, the *EMC Data Manager Oracle Backup Client* guide, the *EMC Data Manager EMC Backint* guide, the *EMC Data Manager Windows NT Oracle Backup Client* guide, the *EMC Data Manager Windows NT SQL Server Backup Client* guide, and the *EMC Data Manager Windows NT Exchange Backup Client* guide.

---

## ACL Support

An Access Control List (ACL) provides an enhanced level of security for UNIX files. ACLs extend the standard UNIX permission settings beyond owner, group, and other. An owner of a file can permit or deny access to specific users and groups.

For a list of platforms that support ACLs, refer to the *EMC Data Manager Software Release Notes*. HP, IBM, DEC, and Sun platforms implement ACLs differently. Refer to the appropriate client documentation for details.

---

## Backing Up Files with ACLs

The backup software retains ACL settings when a file is backed up. During the backup process, backup writes the ACL to the media along with the data. When restored, backup properly restores the data with the same permission settings to the originating or same type client.

IBM clients support file ACLs up to one memory page (approximately 4096 bytes) in size. However, backup does not retain a file's ACL if it exceeds 1024 bytes.

If you attempt to back up a file that has an ACL larger than 1024 bytes, the backup process backs up the file without the ACL data. Only the standard UNIX file permissions are preserved.

This also produces an error message. If this error occurs often, consider adding more user groups to manage file access logically.

---

### Restoring Files with ACLs

When you browse backup catalogs and mark files for restore, ACL settings are not visible in the file listing. However, backup checks the ACL settings and prohibits users who do not have permission to restore the file. A user can have access according to standard UNIX permissions but is prohibited from accessing a file if specified by the ACL. On the other hand, if a user has access to a file via the ACL but does not have standard UNIX permission, the user cannot mark the file for restore.

Due to the way that ACLs were implemented in Solaris 2.5.x, restore of ACLs on *directories* is not supported. (Restore of ACLs on files is supported.) The way Solaris 2.5.x implemented ACLs, the root account cannot change the ACL of a file or a directory of which root is not the owner. The restore software reconstructs directories from their attributes, but since no owner is defined, the restore software is unable to set the ACL.

---

### Cross-Client Restore

The backup software does not support cross-client restore of ACLs. An ACL setting is retained only if you restore the file to the same platform type. For example, if you back up a file with ACL settings from an HP platform you can only restore the file with its original ACL to an HP platform. If you restore the file to another non-HP platform, the file is restored but the ACL is not retained.

## Client ACL Commands

The commands that you use to list and set ACLs differ for each platform. Table 5-1 lists ACL user commands for the HP-UX, IBM AIX, Sun Solaris, and DEC UNIX platforms. Refer to the HP, IBM, Solaris, or DEC UNIX documentation for more information.

Table 5-1

Client ACL Commands

Operating System	Command	Description
HP-UX	<b>chacl</b> (1)	Change ACLs of files
	<b>getaccess</b> (1)	List access rights to files
	<b>lsacl</b> (1)	List ACLs of files
IBM AIX	<b>acledit</b> (1)	Edit an ACL
	<b>aclget</b> (1)	List ACLs of files
	<b>aclput</b> (1)	Set an ACL for a file
Sun Solaris	<b>acl</b> (2) *	Edit an ACL
	<b>aclsort</b> (3) *	Sort an ACL
	<b>getfacl</b> (1)	List ACLs for a file or files
	<b>setfacl</b> (1)	Set an ACL for a file or files
DEC Unix	<b>getacl</b> (1)	List ACLs of files
	<b>setacl</b> (1)	Set an ACL for a file or files

\* System call or library function

## Client Pacing

You can use Client Pacing to control the network bandwidth that the network backup clients of the EDM use. Enabling Pacing frees up computer resources for use by other

applications. Pacing then ensures that the average network utilization over a period of time does not exceed a specified threshold, thus "pacing" resources among applications.

**Note:** Client Pacing is available on all UNIX clients except Auspex and SunOS.

To use of the Client Pacing feature, do the following:

1. Make sure you install (or reinstall) the client after installing the EDM server software.
2. Create a file "pacer.cfg" on the client, in the directory /usr/epoch/EB. This file should have read permission for all users. This is a single line text file with the format:

*Threshold* [*debug\_mode*]

where:

*Threshold* specifies the threshold value in KB/sec. The smallest permissible value for threshold is 100 (KB/s).

*debug\_mode* is optional. Valid values are:

- a. "Verbose," which writes Pacer trace messages to the file pacer.log in /tmp. Verbose is for temporary use; continued use could flood /tmp.
  - b. "quiet," which suppresses Pacer trace messages (the default if no value is specified). For example, to limit the network utilization to 1000 KB/sec, this file should have an entry of 1000.
3. Run your backups as usual.

When the backup process starts on the client, it reads this file. If the file is successfully read and parsed, the pacing feature is enabled, and a message is sent to be logged in the file backups.log, in directory /usr/epoch/EB/log, on the EDM server:

Client Pacing is enabled for this backup. Threshold = 1000

where the threshold value is the one set in `pacer.cfg` on the client. You may easily disable the pacing feature by commenting out the `pacer.cfg` entry using “#” or removing the file `pacer.cfg` from the client.

**Note:** Understand that Client Pacing is done at the expense of the backup throughput. Overall backup performance is, by definition, impacted.

Keep in mind the following important points:

- The threshold value is applied to each backup that may be running for the client. It is NOT a collective threshold for all backups. So, for example, if threshold is set to 1000 KB/sec, and two backups are running concurrently for the same client, each is paced to the order of 1000 KB/sec, and the overall network utilization by all of the backup processes is 2000 KB/sec.
- Once the backup process begins on the client, the `pacer.cfg` file is not read again. Thus, any changes to this file do not affect any backups that are already running.
- The threshold value should be perceived as an approximation. The EDM client attempts to keep the average throughput over a period of time under the threshold value, but the network utilization at any particular instant is not guaranteed to be equal to the threshold.

---

## Server Processes Attributes and Data

When the backup server receives the files from the client, the **ebbbackupd** daemon creates a backup saveset on the server to hold the contents of the backup data stream that it receives from the client system.

The **ebbbackupd** daemons can interleave savesets from different clients on the same piece of backup media, allowing many backups to occur simultaneously.

The allocation and use of backup media is managed by Volume Management through the *trail* concept. A trail is a collection of backup media of the same type, which you expand by allocating new media as needed. You specify the type of media to use for each backup level when you define the trail.

The **ebbackup** program can optionally alternate trails every other day. This enables you to segregate data between separate and identifiable sets of media, which makes it possible to store backups off site.

To restore files, the restore program (called from the Restore window in the EDM interface or from the **ebrestore** command) uses the backup catalogs, which are essentially a snapshot of a client's file names at the time of the backup. The EDM Restore window enables you to browse through the file names at any point in time, and to select individual files or entire filesystems to restore. For more information, refer to the section "How Restore Works" on page 5-16.

---

## Catalog Processing

When a backup is first completed, the raw data for the associated catalog exists on the server, but the catalog daemon (**ebcatalogd**) must process the raw catalog before the restore program can use it. You can have catalog processing performed concurrently with backups, or you can schedule catalog processing for a later time so that this task does not slow down backups.

By default, **ebcatalogd** starts when the system boots. To start and stop processing, add the following commands to the crontab file:

```
/usr/epoch/EB/config/daemon_startup -ebcatalogd  
/usr/epoch/EB/config/daemon_startup -stop
```

Refer to the **ebcatalogd** man page for more information.

---

## Server Database Update

The **ebbbackupd** daemon is responsible for maintaining several server databases, such as the saveset, volume management, and the catalog databases.

---

## Report and Log File Generation

When a backup completes, you can configure and save backup queries to report on through the EDM graphical user interface (GUI), as described earlier. Through this reporting you can gather information about a backup such as its status, total throughput, total size of the backup, and total files that were backed up. (Refer to EDM online help, "Backup Report Overview" for more information.)

At backup completion, the system also generates either a backup completion report to confirm backup operations or a backup failure report. For examples, refer to "Backup Completion Reports" on page 16-29 and "Backup Failure Reports" on page 16-31 for more information.

The **ebbbackup** utility can email success and/or failure reports to the system administrator or to a configured list of users.

The **ebbbackupd** daemon records progress reports in the server log file of each work item as it is backed up.

---

## Backup Completion Reports

Backup completion reports describe successful backups. These reports list the backup template name, the backup start date and time, the name of the backup trail, each client and what was backed up from it, the amount of kilobytes of client data in the backup, any clients that were unavailable for backup, the total number of clients, files, directories, and the backup completion date and time.

The server **eb\_server\_config** installation procedure creates the **mailok** script to which it passes the backup completion information. The script mails the reports to individuals who are responsible for backup operations, and/or writes them to a log.

---

## Backup Failure Reports

Backup failure reports list the backups that require manual intervention to proceed. These reports contain only serious error messages, such as notifications of an interrupted backup or of a backup that could not start. Backup failure reports do not include errors from which EDM Backup automatically recovered.

The server installation procedure creates the **mailerr** script to which it passes the backup failure information. The script can mail the reports to individuals responsible for backup operations, and/or write them to a log.

---

## Backup and Restore Logs

The **ebbackupd** daemon records progress reports in the server log file of each work item as it is backed up. During setup, a *logging level* is specified in each backup schedule template, which controls how much information this file contains. You can modify the default logging level, which is *stats*.

The `/usr/epoch/EB/log` directory on the server contains the following log files:

- `backups.log` — contains an audit trail of backup-related activities listed in chronological order. EDM Backup adds information to this file each time it backs up a template's work items. Selected notifications that appear in this log file also appear in other backup reports. As each log file accumulates information, EDM Backup removes the oldest ten percent of the data after the file reaches a configured maximum size.
- `recoveries.log` — contains an audit trail of restore-related activities that are listed in chronological order. EDM Backup adds information to this file each time it performs a file or database restore for a client.
- A log file for each backup template — contains the backup history for a single backup template. The information in this log varies depending on the logging level you specify in the



configuration database. Thus, you can use the file to view a history of backup-related events for a single template. The log files are named *template\_name.log*.

For more information, refer to "Log Files" on page 16-35.

## How Restore Works

The restore program provides a means of retrieving data from backups, which ensures that lost or damaged data can be quickly replaced.

**Note:** For Symmetrix Connect, see also the *EMC Data Manager Symmetrix Connect User Guide*, the *EMC Data Manager Oracle Backup Client* guide, the *EMC Data Manager EMC Backint* guide, and the various Windows NT client guides.

Table 5-2 gives you an understanding of the valid restore paths for the three types of backup paths.

Table 5-2

Backup versus Restore Path

Backup Path	Restore Path		
	Network	Symmetrix Path	Symmetrix Connect
Network	Yes	Yes	No
Symmetrix Path	Yes	Yes	No
Symmetrix Connect	Yes <sup>1</sup>	Yes <sup>1</sup>	Yes

1. You cannot restore Symmetrix Connect data via network or Symmetrix Path if your Symmetrix Connect backup is a striped LVM configuration. You also cannot restore a Windows NT Symmetrix Connect backup via network or Symmetrix Path.

The restore program offers self-service file retrieval. You can configure the backup server so that users on UNIX client systems can perform file restore without the aid of a system administrator.

The EDM Backup server tracks user names and the clients from which these users have permission to restore files. This provides system security by enabling access control on a per client system basis. The software also enforces UNIX file permissions and Access Control Lists, if supported by the client. Users can restore only those files for which they have access permission.

**Note:** Refer to your *EMC Data Manager Software Release Notes* for a current list of client platforms that support ACLs.



Click this icon in the Main window of the EDM GUI to open the Restore window, or enter the command **edmrestore** at the CLI.

Use the CLI command **edmrestore** to open the Restore window on clients. The EDM makes the connection and opens the Restore window. Alternatively, users can issue the command **ebcrecover** to restore files.

**Note:** The **ebcrecover** command on the client calls **ebcrecover** on the EDM, but as of EDM 4.5.0, **ebcrecover** is just a symbolic link to the actual restore program, **ebrestore**.

(Refer to the **edmrestore**, **edmcrestore**, and **ebcrecover** man pages for more information about these commands.)


The restore program (**ebrestore**) runs on the server and enables you to extract data from backup media. Restore works with the help of on-line catalogs that keep track of all files that are saved on each backup volume.

**Note:** Try to avoid restoring a work item that is currently being backed up.

Figure 5-2 shows an overview of the file restore process:

Figure 5-2

## The File Restore Process



1. A user on a client system issues the **edmcrestore** command to connect to the server and start up the restore window.

During one restore session, a user can restore files from different backup savesets. This causes the restore program to combine catalogs to reproduce the state of the filesystem at the time of the request.

2. A user searches for and marks files to restore. The server then locates the backup savesets that contain the specified files.

**Note:** When restoring offline database backups, you must select all stripes from the same backup date. You cannot restore one stripe, for example, from one backup date and another from a previous date.

3. The backup program scans the saveset records, which identify the backup catalogs, and locates the media that stores the backup data.

4. The **restore** program connects to the client using the client's native connection method (or direct connection in the case of a local client).

In either case, an EMC-provided restore shell script is invoked on the client system. The backup daemon provides the client script with a specification of the files to be restored.

5. The **restore** program reads the files being restored from the server's storage devices, packages the files in an extended **cpio** format, and sends them to the client system.
6. The client system writes the files to its disk.

If the files are being restored to the original client, the restored files overwrite any files that already exist on that client (unless the user specifies an option to not overwrite existing files). If the files are being restored to a different client, the server restore the files to the requested location.

7. Information about the restored files is written to the `recoveries.log` file on the client and on the server.

The `recoveries.log` includes the name of the user who is performing the restore operation, the total amount of data restored, and the date and time of the restore.

---

## Media Management

The server's volume management software manages all media for backup and HSM (for network backups only). The software controls the library unit robotics that automatically insert media into and remove them from drives when an application requires their use.

For an understanding of media management, refer to Chapter 7 "Basic Volume Management Concepts" and Chapter 8 "How Volume Management Works."

---

## Expiring Backups

As time passes, your site will have many volumes of backup data. If the site no longer needs to keep older backup data, it can specify that backups that are older than a certain age be automatically expired. When every backup with data on a volume has expired, the media is made available again. The volumes that are made available via expirations are made available only for their current trail. You may relabel these volumes to remove this restriction.

---

## Deallocating Baseline Volumes

In HSM systems, baseline volumes do not expire directly but become deallocated (that is, available for reuse) when three conditions are met:

- The volume is no longer the current volume for the trail.
- The volume is not referenced by any baseline-relative backup.
- No block on the volume is in use; that is, no data is actively referenced by the current file system (as when a baseline ID in a file's metadata points to the volume).

---

## 6 Database Backup and Restore

Various options are available for database backup and restore. Backups can be either over the network or over SCSI or Fibre Channel cabling from a Symmetrix storage system. This chapter describes the several ways to back up databases.

**Note:** Prior versions of EMC Data Manager included an "offline" backup feature for Oracle, Sybase, and Informix databases — no option required. EDM no longer includes this feature. See "EDM's Legacy "Offline" Database Backup Feature" on page 6-17.

This chapter includes the following topics:

- Varieties of Database Backup
- Various Database Backup Clients
- Database Network Backup Overview
- EDM Symmetrix Path Overview
- EDM Symmetrix Connect Overview
- EDM's Legacy "Offline" Database Backup Feature

---

## Varieties of Database Backup

EDM supports database backup and restore over networks or through attached Symmetrix storage units. Database support is provided by optional backup clients and by two Symmetrix options, EDM Symmetrix Path and EDM Symmetrix Connect.

The database backup clients provide network backup for the various major databases on numerous system platforms. In addition, all of the clients work with EDM Symmetrix Path (on the most popular system platforms). Finally, some of the clients work with EDM Symmetrix Connect as applicable, but EDM Symmetrix Connect also provides specialized backup interfaces in some cases to support database backup.

Whichever backup methodology used, the EDM centrally manages backup processing and media operation. Provided that adequate tape drives are available, EDM can run Symmetrix-related and network backups at the same time.

EDM's network backup is an effective solution for data centers with many small to medium-sized database and file servers. Network backups and restores are accomplished through ATM, FDDI, Fast Ethernet, or Token Ring network connection(s). The network backup function scales well, performs multiple backups at once, and centralizes backup and media management.

The EDM Symmetrix Path option offers channel-speed backup of large databases (and filesystems) over a data path through Symmetrix storage units to the EDM. Backups and restores are done over Fibre Channel or SCSI cables connected between the database host, the Symmetrix, and the EDM.

The EDM Symmetrix Connect option specifically addresses the backup needs of very large databases. Backups and restores are done over Fibre Channel or SCSI cables connected directly between the Symmetrix and the EDM, providing a direct backup of the database (or its mirrored-copy) on Symmetrix storage units to the EDM.

---

## Database Backup Clients

EDM's database backup clients support various databases from major vendors, specifically: Oracle<sup>™</sup>, SAP R/3<sup>™</sup> System (Oracle database), Sybase®, Informix<sup>™</sup>, Microsoft SQL<sup>™</sup> Server, and Exchange. Unix and Windows NT platforms are supported. (See the *EMC Data Manager Software Release Notes* for a current list of supported client platforms and operating system versions.)

For the most part, these database backup clients provide backup and restore of database systems through interactions with the standard backup utilities provided by the database vendor for use with their databases. For example, the EDM Oracle Backup Client supports backups by Oracle7's Enterprise Backup Utility (EBU) and Oracle8's Recovery Manager (RMAN).

Depending on the database client, backup and restore operations are initiated from the database system, from the EDM, or from either system.

---

## EDM Symmetrix Path

The EDM Symmetrix Path option is available for use with all the database backup clients. See the *EMC Data Manager Software Release Notes* for a current list of database versions and operating system versions supported for EDM Symmetrix Path.

Instead of streaming backup data over the network from the database host to the EDM, EDM Symmetrix Path streams the data over the Fibre Channel or SCSI cables between the database host, the Symmetrix, and the EDM. (The network is still used for control communication between the EDM and the client.) Once configured for EDM Symmetrix Path, a client's backups and restores can be switched between EDM Symmetrix Path and the network with a simple reconfiguration.

In supporting EDM Symmetrix Path, the database clients interact with the database backup utilities in the same way as they do for network backup. For example, the EDM Oracle Backup Client supports EBU and RMAN backups over Symmetrix Path.



---

## EDM Symmetrix Connect

The EDM Symmetrix Connect option supports backup of a few key database systems, (Oracle, Microsoft SQL Server™, and Exchange) residing on Symmetrix storage units. It also supports backups of filesystems.

With EDM Symmetrix Connect backups, data on the Symmetrix (either the database or a mirrored-copy) is streamed over the cables directly from the Symmetrix to the EDM. The network is still used for control communication between the EDM and the client and for backup and restores of archived redo logs and control files.

EDM Symmetrix Connect backs up data for a few key database applications and filesystems. Two of its application interfaces are standard; the others are EDM-specific interfaces.

EDM Symmetrix Connect backs up UNIX Oracle databases through two standard interfaces:

- RMAN Proxy Copy (for some UNIX platforms), enabling backups and restores through Recovery Manager (RMAN)
- EMC Backint, SAP-certified interface (for some UNIX platforms), enabling backups and restores through SAPDBA

The RMAN Proxy Copy application supports Oracle's RMAN utility's Proxy Copy backups in conjunction with the EDM Oracle Backup Client. For RMAN Proxy Copy, backup operations can be initiated from the database system or the EDM; restores can be initiated from the database system only.

The EMC Backint application supports the SAP R/3 System's SAPDBA backups of Oracle databases, using the SAP R/3 System's standard BACKINT interface in conjunction with the EMC Backint client. For EMC Backint, backup and restore operations are initiated from the database system.

In addition, EDM Symmetrix Connect has two EDM-specific interfaces to Oracle:

- EDM Oracle Application Interface (for several UNIX platforms), enabling backups and restores through EDM
- EDM interfaces for Windows NT systems for Oracle, Microsoft SQL Server, and Exchange

With the EDM Oracle Application Interface, EDM Symmetrix Connect is tailored to work with directly with Oracle databases on the following client platforms: Compaq, HP, IBM, Sequent, and Sun. Rather than interacting with the Oracle backup utilities, the EDM Symmetrix Connect software provides its own interface to the Oracle database.

Similarly, EDM Symmetrix Connect's Windows NT interfaces support backups of Oracle, Microsoft SQL Server™, Microsoft Exchange™ (as well as filesystems on Windows NT). See the *EMC Data Manager Software Release Notes* for a current list of supported client platforms.

For the EDM-specific interfaces, backup and restore operations are initiated from the EDM.

---

## Various Database Backup Clients

EDM supports various, distinct database clients that are packaged separately as options.

---

---

### Oracle Backup Client

The EDM Oracle Backup Client, which interacts with the Oracle backup utilities, is available for many UNIX client platforms and Windows NT. It supports network, EDM Symmetrix Path, and EDM Symmetrix Connect (RMAN Proxy Copy) backup. This option is installed and configured through the EDM graphical user interface (GUI). See *EMC Data Manager Oracle Backup Client* for installation and configuration instructions for UNIX platforms. For Windows NT, see the *EMC Data Manager Windows NT Oracle Backup Client* manual. Also see the online Help in the GUI.

---

---

### EMC Backint Client for SAP R/3 Oracle Databases

The EMC Backint client, which is a software interface to the SAP R/3 System for backing up its Oracle database, is available for many UNIX client platforms and for Windows NT. It supports network, EDM Symmetrix Path, and EDM Symmetrix Connect backup. This option is installed and configured through the EDM GUI. EMC Backint has its own client manual with specific installation and configuration instructions. See the *EMC Backint* manual.

---

---

### Other UNIX Database Backup Clients

Sybase and Informix databases each have corresponding database clients, with individual client manuals with specific installation and configuration instructions, which are available for most client platforms. They support network backup and EDM Symmetrix Path. These options are installed and configured through the EDM GUI. See the *EMC Data Manager Sybase Backup Client* and *EMC Data Manager Informix Backup Client* manuals.

---

---

## Microsoft Database Backup Clients

Microsoft SQL Server and Microsoft Exchange each have corresponding database clients, with individual client manuals with specific installation and configuration instructions. They support network, EDM Symmetrix Path, and EDM Symmetrix Connect backup. These options are installed and configured through the EDM GUI. See the *EMC Data Manager Windows NT SQL Server Backup Client* and *EMC Data Manager Windows NT Exchange Backup Client* manuals.

---

## Database Network Backup Overview

Database backups can be initiated from within the database management system or on the EDM. In the first case, the database administrator starts the backups from within the DBMS's own backup utility. In the second case, the EDM's scheduling function starts a process that, in turn, automatically starts the backups within the DBMS's own backup/restore utility.

Table 6-1 lists the backup and restore utilities for each database system.

---

Table 6-1

---

**Network Database Backup and Restore Utilities**

Database System	Database's Backup/Restore Utility
Oracle8	Recovery Manager (RMAN)
Oracle7	Enterprise Backup Utility EBU (obackup)
Oracle under SAP R/3	SAPDBA (BRBACKUP, BRARCHIVE, BRESTORE)
Sybase	dump/load
Informix	On-BAR

---

Once started, the DBMS's backup utility scans the database for data to back up and passes the data to the backup client software. Then the backup client software streams the data to the EDM over the network.

---

## Multiple Streams

To speed up backup processing, you can create multiple streams of backup data from the database (for example, six). You can also decide to write the backups to multiple tapes (for example, two). The difference in the two example numbers represents a consideration of the generally slower speeds of magnetic disks (which hold the database) as compared with tape drives (which write the backups to tape).

**Note:** The various implementations of streaming for each database client are described in their respective manuals (or release notes, as applicable).

---

## Server-side Processing

The server software manages the writing of data to tape storage media and provides online catalogs, located on EDM's own disks. Separate paths for data flow and for control (catalog information), allows the data to take a more direct path to the backup media.

The catalogs enable restores of your data from the backup media. (In the case of the Backint SAP R/3 client, the client software also creates catalogs at a meaningful granularity for the database system and stores them locally on the client.)

The server software also manages the operation of robotic library units and provides overall volume life-cycle management. Also, the software enables automated backup expiration.

---

## Restores

The client software receives restore requests from the database's restore utility at a database, tablespace, or data file granularity, as appropriate, and sends the requests to the server software.

The server software retrieves the data from the backup media and sends it to the client software, which passes it on to the database's own restore utility.

---

## User Interfaces

The backup clients can be installed through the EDM Backup Install Wizard and configured through the EDM Backup Configuration Wizard. They can be reconfigured through either the EDM Backup Configuration Wizard or the EDM Backup Configuration window. Some command-line procedures might also be required.

See the appropriate EMC manual for information on installing clients, configuring backups, and restoring data.

**Note:** Although the Restore window might display work items for the following database products, it does not support restores for: Oracle Backup Client, Sybase Backup Client, Informix Backup Client, and EMC Backint. Restores are accomplished through each database's restore utility.

---

## Configuring Backups to an Alternate Network

Multiple-networked backup, meaning that if the *client* machine has multiple interfaces, the client can send backup data through those multiple interfaces.

You can specify which interface to use by Work Group/Schedule, using the Use Client Name parameter in the Work Item tab's Work Item Options window (in the Generic tab). This parameter corresponds to the *connection* *via* parameter of the "listener" work items in the eb.cfg file.

**Note:** A separate issue is the use of multiple network interfaces on the *EDM*. For database backup clients only, the client must have a valid EB server hostname in its `/usr/epoch/EB_DB/ebci.conf` file.

---

## Database Pre-Discovery

During initial client configuration from the Backup Configuration Wizard, the client machine is scanned for the presence of databases, in a process called *pre-discovery*.

Pre-discovery reveals just two pieces of information:

- Database type (that is, Oracle, Sybase, Informix)
- Database name

No special requirements exist for Oracle databases to be pre-discovered: the Backup Configuration Wizard presents the databases listed in the `/etc/oratab` file as the databases that exist on the system. However, for Sybase and Informix databases, there are assumptions as to where the database software was installed, as described in Table 6-2.

**Table 6-2 Requirements for Database Pre-Discovery**

Database System	Requirement for Pre-Discovery
Oracle	<p>None. The Backup Configuration Wizard presents the databases listed in the <code>/etc/oratab</code> file as the databases that exist on the system.</p> <p><b>Note:</b> If the file lists a database that actually does not exist, it is shown anyway, as pre-discovery is not able to ascertain whether in fact the database does not actually exist or if it is just off line.</p>
Sybase SQL Server	<p>One of the following on the Sybase machine:</p> <ul style="list-style-type: none"><li>• A link, <code>"/SYBASE"</code>, to the directory in which the Sybase database software is installed.</li><li>• The environment variable <code>"SYBASE"</code> for the root environment pointing to the directory in which the Sybase database software is installed.</li><li>• A UNIX account named <code>"sybase"</code> in whose home directory (as specified in <code>/etc/passwd</code>) the Sybase database software is installed.</li></ul>
Informix	<p>Either of the following on the Informix machine:</p> <ul style="list-style-type: none"><li>• The environment variable <code>"INFORMIXDIR"</code> in the root environment pointing to the directory in which the Informix database software is installed.</li><li>• A UNIX account named <code>"informix"</code> in whose home directory (as specified in <code>/etc/passwd</code>) the Informix database software is installed.</li></ul>



---

## EDM Symmetrix Path Overview

The EDM Symmetrix Path feature currently works with the EDM Oracle Backup Client, the EMC Backint client, Sybase client, Microsoft NT SQL Server client, Microsoft Exchange client, and Informix client.

EDM Symmetrix Path enables an EDM to back up your database (and filesystems) through direct SCSI connections to a Symmetrix, rather than over a local area network.

With this methodology, the Symmetrix itself acts as the network. A few small devices are designated as transport paths for configuration purposes, while the actual data transport is generally handled by the cache on the Symmetrix corresponding to these devices.

**Note:** As they are dedicated to transport, these devices are unavailable for use as storage devices.

These Symmetrix Transport Groups (also referred to as ST groups or STGs) are mapped to the hosts (clients and the EDM). Various possible device configuration and host mapping combinations provide different performance characteristics and degrees of flexibility.

**Note:** The database can be located on the Symmetrix on volumes that are accessible to the database host. But unlike with the EDM Symmetrix Connect methodology, the database volumes are not made accessible to the EDM, nor are mirror-images of the database volumes.

Each of the various possible device configuration and host mapping combinations provide different performance characteristics and degrees of flexibility. See the *EMC Data Manager Symmetrix Path User Guide* for more information.

---

## **EDM Symmetrix Connect Overview**

With EDM Symmetrix Connect, databases are located on one or more Symmetrix systems in various configurations that are visible to the EDM. The backup data is sent over SCSI or Fibre Channel cabling that directly connects the EDM to the Symmetrix storage. The backup is usually taken from a mirrored copy of the database on a Symmetrix (three different mirrored-volume configurations are available), but a non-mirrored configuration, which backs up the database itself, is also supported.

---

## **Applications**

EDM Symmetrix Connect's EDM Oracle Application Interface supports backup of Oracle databases for UNIX (Sun, IBM, HP, Sequent, and/or Compaq clients). The Filesystem Application of EDM Symmetrix Connect adds the capability to back up selected filesystems on UNIX systems.

Other EDM Symmetrix Connect applications support Oracle backups through RMAN Proxy Copy and through EMC Backint (for SAP R/3 System Oracle databases). Both online and offline database backups can be performed.

EDM Symmetrix Connect's Windows NT applications support Oracle, Microsoft SQL Server, and Microsoft Exchange (as well as Windows NT filesystems).

---

## **User Interfaces**

On the EDM, the EDM Backup Configuration Wizard is used for client configuration of EDM Symmetrix Connect.

The Backup Activity Wizard or EDM's command-line interface is used for EDM-initiated backup.

Restores of EDM Symmetrix Connect backups taken through the EDM Oracle Application Interface are EDM-initiated using the command-line. Restores of Oracle RMAN Proxy Copy backups

are always performed on the client, through RMAN. Restores of EMC Backint backups are always performed on the client, through SAPDBA or BRRESTORE.

The EDM GUI's Library Manager window is available for media management operations.

---

## Documentation

See the *EMC Data Manager Symmetrix Connect User Guide* for detailed information on how to perform each operation for the EDM Oracle Application Interface and for the filesystem applications.

See the *EMC Data Manager Oracle Client* guide for RMAN Proxy Copy, the *EMC Data Manager Backint* guide for EMC Backint, and the *EMC Data Manager Windows NT Oracle Backup Client* guide for NT Oracle.

See the other Windows NT client guides corresponding with the EDM Symmetrix Connect backups of your Windows NT system.

---

## Raw Device Backups

For the most part, EDM Symmetrix Connect performs its database backups at the physical disk, that is, raw device level. Raw device backups have the advantage of fast backup performance. Their limitation is that they do not offer the granularity of file-by-file (logical) backups and restores.

If your database files are built on raw devices, then the raw device backups are perfectly suited.

If your database files are in filesystems, logical filesystem backups are available if your database server (the backup client) is running the same operating system as the EDM.





Otherwise, raw device backups of database files in filesystems will give you fast backup performance but also a lack of file-by-file granularity in backups and restores.

## Configurations

Databases can be located on one or more Symmetrix systems in various configurations. EMC offers a direct connect backup and restore solution in a combination with four Symmetrix volume mirroring configurations as shown in Table 6-3.

**Table 6-3**

**EDM Symmetrix Connect Mirroring Configurations**

Backup Solution	Symmetrix <sup>1</sup> Configuration	Description
Symmetrix SRDF Connection		Backup of "target" ("R2") volumes on a second, connected Symmetrix (b), which mirror the "source" ("R1") volumes on the first Symmetrix (a), which contain the database
TimeFinder		Backup of "Business Continuance Volumes" ("BCVs") (b), which mirror the "standard" ("STD") volumes (a), which contain the database, all of which are on one Symmetrix
Remote BCV		Backup of "BCVs" (c) on second Symmetrix, which mirror "target" ("R2") volumes on the second Symmetrix (b) that mirror the "source" ("R1") volumes on the first Symmetrix (a), which contain the database
Non-mirrored		Backup of non-mirrored volumes (a) containing the database

1. For all EDM Symmetrix Connect solutions, the Symmetrix models are 3xxx/5xxx ESP model systems

## Mirrored Configurations

The first three configurations listed in Table 6-3 take advantage of volume mirroring, which is when one (or more) exact copies of the database's disks are maintained simultaneously.

One mirroring capability is between two physical Symmetrix systems (the Symmetrix SRDF configuration). The other mirroring capability is between disks on a single Symmetrix (the TimeFinder configuration). The Remote BCV configuration uses both mirroring capabilities together.

To perform the backups, the EDM logically discontinues the active mirroring between the database disks (or the target "R2" volumes, in the case of Remote BCV) and the mirrored copies. The database host can continue to function normally using the primary Symmetrix disks while the EDM backs up the (now static) mirror-copy of the database (or tablespaces) rather than the database itself. At the time the mirrors were split, the two volume sets were exact copies of each other.

Backups can be online or offline. In a mirrored configuration, when performing online backups, the database is put into Oracle's online backup mode only briefly, just for the time the mirrors are split. (If any inconsistencies occur while the mirrors are split, they can be resolved by the Oracle software at restore time.) Therefore, for these mirrored configurations, the database host (client) is not impacted by the online backup operation.

When performing offline backups with the EDM Oracle Application Interface, the database or tablespace remains offline only briefly, just for the time the mirrors are split. However:

- For EMC Backint, the database or tablespace remains offline for the entire duration of the backup.
- For RMAN Proxy Copy, the database, tablespace, or datafile remains offline for the entire duration of the backup unless a post-mirror-split script is employed to bring the database, tablespace, or datafile online after mirrors have been split.

---

### **Non-Mirrored Configuration**

In the Symmetrix non-mirrored configuration, the actual client's database is backed up, not a mirrored copy of it. To use this feature, the Oracle data files cannot reside in a filesystem. The data files must be in raw partitions.

The database can be backed up either online or offline. When performing online backups in a non-mirrored configuration, the database is put into Oracle's online backup mode for the entire duration of the backup.

**Note:** The extended time for online backup mode may be an issue for certain customers since Oracle overhead can be significant, if major updating is being performed during online backup mode. The effects of online backup mode overhead can be reduced by configuring backups, such that fewer tablespaces are requested per backup run.

When performing offline backups, the database or tablespace remains offline during the entire duration of the backup.

---

### **EDM's Legacy "Offline" Database Backup Feature**

EMC Data Manager formerly included an "offline" database network backup feature for Oracle, Sybase, and Informix databases, for which no option was required.

Support for performing backups with this feature has been removed.

However, EDM does continue to support restores for any backups taken with this functionality using prior versions.

Also, EDM continues to support backups of these databases on most of the same client platforms through the use of EDM's Oracle, Sybase, or Informix backup clients.



---

---

## **7 Basic Volume Management Concepts**

Volume management software manages and controls access to attached library units and all removable media that EDM Backup and the optional HSM software use.

This chapter describes the components that comprise the volume management software.

The topics in this chapter include:

- Volume Management Overview
- EDM Library Unit Manager
- Volume Manager
- Library Managers




## Volume Management Overview

The volume management software provides removable media management services to Backup and optional HSM software. The software includes the following components:

- EDM Library Manager (GUI and CLI)
- Volume Manager
- Device-specific Library Managers

Figure 7-1

Volume Management Software



---

## EDM Library Unit Manager

The EDM Library Unit Manager, which resides on the EDM Backup server, is your interface to volume management functions. For example, you use the EDM Library Unit Manager to:

- monitor volume activity on the EDM server
- inject and eject volumes from the library unit
- label volumes for use
- locate available volumes
- check for outstanding media requests
- initiate partial or complete library unit inventories

The interface interacts with the underlying software components, which the following sections describe.

---

---

## Volume Manager

The Volume Manager manages:

- information about all volumes
- volume life cycle
- volume requests that applications make (for example, backup and media duplication)

The Volume Manager maintains a group of files on the server in the directory `/usr/epoch/etc/vm`. This directory contains a configuration file (`vm.cfg`), the volume catalog (`volumes`), template catalog (`templates`), log file (`clog`), and other administrative files.

**Note:** With the exception of `vm.cfg`, volume management uses all of the files in this directory internally. You should not edit any of these files.

---

## Volume Catalog

Volume management identifies all volumes by a unique electronic volume label on the media. The Volume Manager keeps track of all volume information in the volume catalog. The catalog contains entries for each volume including the volume ID, volume sequence number, physical location (by Library Manager), volume state, optional barcode ID, and usage count.

The Volume Manager updates the catalog as operations and events occur, such as when:

- new media enters a library unit
- a volume is allocated to an application
- volumes in a library unit are inventoried
- a volume is ejected from the library unit
- a volume is imported from another server
- a volume is moved from offline to offline

The volume catalog is an integral part of volume management software and necessary for rebuilding an EDM system. Therefore, the catalog as well as all volume management system files are backed up as part of the default server work group.

---

## Volume Life Cycle


A volume's life cycle (as illustrated in Figure 6-2) begins when you label a new tape or optical disk. The labeling process writes a unique electronic label to the media. A labeled tape, such as DLT, holds one volume. Two-sided media, such as EO disks, contain two volumes, one volume for each media side.

When media is loaded into a library unit, its volume label is read before information for the volume becomes visible in the EDM Library Unit Manager. A volume is identified by its

sequence number or one of the following states: uncataloged, unlabeled, unverified, foreign, expired, or erasing (EO disks only). These states are described below.

Figure 7-2

Volume Life Cycle



### Uncataloged

An *uncataloged* volume indicates that the volume has a valid volume label but is missing from the volume catalog. This can happen if the volume was labeled on another server or if the volume catalog was lost as a result of a server disk crash (see Chapter 20 “Recovering a Server from a Disk Failure”).

If you import the volume, the Volume Manager adds an entry to the receiving server’s volume catalog. The volume retains its volume ID and volume sequence number that the Volume Manager of the originating server assigns.

Volume management does not allow duplicate volume sequence numbers of the same media type. If you attempt to import a volume that has a sequence number that already exists, you are prompted to either delete the existing volume or cancel the import operation. (Refer to “Duplicate Volume

Sequence Numbers" on page 8-26 for instructions on how to override this restriction.) Volumes that are contained on different media types (for example, EO and DLT) may have the same sequence number.

---

**Unlabeled**

An *unlabeled* volume means the volume's label area is blank. This is the state of new media. You can either label the volume or leave it unlabeled until an application makes a request for a new piece of media.

When you insert an unlabeled volume into a library unit, the Library Manager assigns it a slot number, moves the volume into the slot, and notifies the Volume Manager to add an entry for the unlabeled volume to the volume catalog.

---

**Foreign**

A *foreign* volume is any previously-used media from a non-EDM system. A **tar** tape is an example.

To reuse a foreign volume, you must first label it. When a foreign volume is labeled the Volume Manager adds it to the volume catalog and changes its state to available. If the media is an EO disk, the data on the disk is erased before it is labeled.

The EDM Library Unit Manager also allows you to mount foreign volumes for the purpose of reading or extracting data. You can do this by manually (or force) mounting the volume into a drive using the Force Mount button in the Utilities tab of the EDM Library Unit Manager window.

**Note:** You must dismount the volume manually to avoid preventing other processes from using the drive.

---

**Expired**

Tape media expires when it reaches a pre-set maximum usage. After backup deallocates a volume, the Volume Manager checks the usage count and expires the volume after the maximum is reached. (See Figure 7-3 on page 7-9.) When a volume expires, it cannot be mounted for data access.

---

## Erasing

Erasable optical disks have an erasing state that occurs at different stages. EO disks enter the erasing state before a foreign EO is labeled and after HSM deallocates the volume. (See Figure 7-4 on page 7-10.)

---

## Unverified Volume

An *unverified* volume means that the Library Manager is unable to recognize the label contents on the volume. An unverified volume can result for one of many reasons: the volume was just injected and did not yet complete the initial label read; the media is incompatible, an error occurred while the label was being read; a hardware problem occurred; or a user placed several volumes into a library unit (LU) through the mass load door and then ran a barcode-only inventory.

Generally you should inventory an unverified volume so that its label is read properly. For a library unit that supports barcodes, perform a barcode and label inventory; for a non-barcode LU, perform a label inventory.

**Note:** A barcode and label inventory is recommended.

If a drive becomes dirty while the volume is in the drive and its label is being read, the Library Manager automatically dismounts the volume, disables the drive and marks it as dirty, and tries to read the volume's label in another drive. If this second drive also becomes dirty, the Library Manager dismounts the volume, places it back in its slot, and marks the volume as unverified and offline. The Library Manager then disables this second drive and marks it as dirty, and places it back into the LU.

If the LU contains a cleaner cartridge, the drive is cleaned automatically; if no cleaner resides in the LU, injecting a cleaner starts an automatic cleaning. You must then verify the unverified volume either through a mount request or an inventory of that slot (barcode/label or label inventory).

**Note:** If the volume remains unverified after an inventory, remove it from the library unit.

---

## How Volumes are Allocated

A volume becomes available for allocation after it is labeled. Labeling a volume requires that you choose a volume template. Volume templates enable you to specify whether a volume should be made available to any application (that is, Backup or HSM), to any trail, or to a specific trail.

Several volume templates are available:

- Unrestricted (*media\_type*) — any application
- EBprelabel — backup only
- HSMprelabelEO — HSM only
- Restricted to *trail\_name* — specified trail name only

The templates also contain attributes that the volume inherits such as: a unique volume ID, trail name, media type, maximum usage count, and media size information.

---

## When a Volume is Allocated

Volume allocation begins with a request from backup, media duplication, or HSM. The Volume Manager locates an available volume based on the application's request, provides a volume ID, gives the application access to the volume, and changes the volume state to *Allocated*.

Volume usage is based on the number of times that a volume transitions from *Available* to *Allocated*. Each time the volume is allocated, the Volume Manager increments the media's use count by one. When a volume reaches its maximum usage, the volume is expired.

The figures on the following pages illustrate the volume life cycle of two media types. Figure 7-3 illustrates the life cycle of tape, Figure 7-4 on page 7-10 illustrates the life cycle of erasable optical media, and Figure 7-5 on page 7-11 illustrates the life cycle of WORM (Write One Read Many) optical media.

Figure 7-3

## Tape States






Figure 7-4

EO Volume States

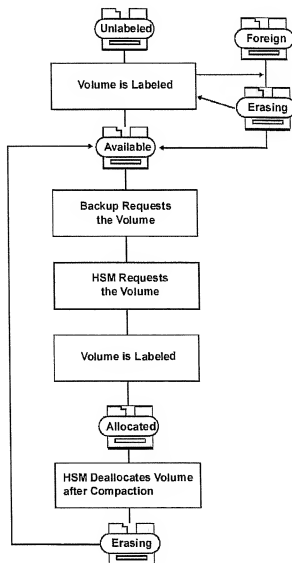



Figure 7-5

## WORM Volume States



---

## Library Managers

Device-specific Library Managers control library unit operations:

- library unit inventories
- drive preemption
- robot movement for mounting and dismounting media
- injecting or ejecting media

Offline and offsite Library Managers hold information about volumes in offline and offsite locations. Library Managers are supplied for various types of library units.

---

## Library Manager Configuration

You use the **lmconfig** utility to configure a Library Manager for each library unit that is attached to the server. For each Library Manager that you configure, **lmconfig** sets up a subdirectory in `/usr/epoch/etc/lm` to include the configuration file and other internal files that the Library Manager uses. Each subdirectory name is based on the vendor name and model number.

Within the Library Manager's subdirectory, a configuration file (`lm.cfg`) defines features and functionality for the device. For example, if a library unit is equipped with a barcode scanner, the configuration file enables barcode support.

After a library unit is configured, an icon for the library unit and each internal drive appears in the Library Units and Drives area of the Library Unit Manager window (only if volume management is running).

---

## Robotic Library Units

A Library Manager controls the library unit's robot (the mechanism that moves cartridges within a library unit), internal tape drives, and media inlet (if present).

Each Library Manager maintains a per-library unit inventory of volumes in the file `valid.dat` within its subdirectory. When a Library Manager is started for the very first time, it takes a complete inventory of the library unit's contents and creates the `valid.dat` file. Once the file is created, the Library Manager reads `valid.dat` to initialize the library unit, which eliminates a complete inventory each time the system is started.

The inventory list includes a volume ID, barcode label (if supported by the library unit), slot number, and drive location for each volume. As volumes move from one location to another, the Library Manager updates the inventory list and notifies the Volume Manager of any changes.

The Library Manager also controls drive scheduling and drive selection. When the Volume Manager makes a request for an operation (for example, a mount request) the Library Manager adds the request to a prioritized work queue. When a drive becomes available, the Library Manager services the next work item with the highest priority.

---

## Offline and Offsite Library Managers

Offline and offsite Library Managers enable you to track the location of volumes that are outside of a physical library unit.

Offline represents volumes that are ejected from a library unit and stored in a nearby area, usually somewhere on site. The offsite Library Manager holds information about volumes that you physically move to a location beyond the building's boundaries, such as an offsite archival location. Only volumes that have a volume label or barcode label can enter the offline or offsite Library Manager. An unlabeled volume with no barcode is deleted from the volume catalog when it is ejected from the LU.

The EDM Library Unit Manager window displays an icon for offline\_0 and offsite\_0 Library Managers. From this window, you can view the volumes that are contained in both the offline and offsite Library Managers. You can also eject volumes into either the offline or offsite Library Manager using the Eject tab in the Library Unit Manager window. (Refer to EDM Online Help for instructions.) The Utilities tab in this window provides a text field that enables you to record the volume's actual offline or offsite location.

**Note:** Moving backup media offsite before its rotation period ends causes the backup that would use that media to fail. You can avoid a failed backup by using new media for the next backup. You configure the use of new media in the Backup Configuration window of the EDM GUI. Select Advanced Options in the Schedule Tab. In the Schedule Options window that appears, select Use New Media When Current (backup media) Is Offsite.

The offline\_0 and offsite\_0 configuration files specify the action to be taken when a mount request is received for a volume that is offline or offsite.

---

### Ejecting a Volume

When you eject a volume from a library unit by using the Eject tab in the Library Unit Manager window, the default destination is offline. (Volume Manager determines where to put the volume by the value of LM\_EJECT\_DEST that is set in the vm.cfg file.)

The Eject tab also enables you to eject volumes to offsite. Knowing whether a volume is located offline or offsite is helpful when you need to locate a volume for an application such as a restore. If the volume is not in any library unit, the Media Request window specifies whether the volume is offline or offsite.